



Red Hat Developer Toolset 2.x User Guide

Installing and Using Red Hat Developer Toolset

Jaromír Hradílek

Jacquelynn East

Matt Newsome

Red Hat Developer Toolset 2.x User Guide

Installing and Using Red Hat Developer Toolset

Jaromír Hradílek
Red Hat Engineering Content Services
jhradilek@redhat.com

Jacquelynn East
Red Hat Engineering Content Services
jeast@redhat.com

Matt Newsome
Red Hat Software Engineering
mnewsome@redhat.com

Legal Notice

Copyright © 2013, 2014 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](#). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux® is the registered trademark of Linus Torvalds in the United States and other countries.

Java® is a registered trademark of Oracle and/or its affiliates.

XFS® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack® Word Mark and OpenStack Logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Red Hat Developer Toolset is a Red Hat offering for developers on the Red Hat Enterprise Linux platform. The Red Hat Developer Toolset User Guide provides an overview of this product, explains how to invoke and use the Red Hat Developer Toolset versions of the tools, and links to resources with more in-depth information.

Table of Contents

| | |
|--|-----------|
| Preface | 10 |
| 1. Document Conventions | 10 |
| 1.1. Typographic Conventions | 10 |
| 1.2. Pull-quote Conventions | 11 |
| 1.3. Notes and Warnings | 12 |
| 2. Getting Help and Giving Feedback | 12 |
| 2.1. Do You Need Help? | 12 |
| 2.2. We Need Feedback! | 13 |
| Part I. Introduction | 14 |
| Chapter 1. Red Hat Developer Toolset | 15 |
| 1.1. About Red Hat Developer Toolset | 15 |
| 1.2. Main Features | 17 |
| 1.3. Compatibility | 18 |
| 1.4. Getting Access to Red Hat Developer Toolset | 19 |
| 1.4.1. Using RHN Classic | 19 |
| 1.4.2. Using Red Hat Subscription Management | 20 |
| 1.5. Installing Red Hat Developer Toolset | 21 |
| 1.5.1. Installing All Available Components | 22 |
| 1.5.2. Installing Individual Package Groups | 22 |
| 1.5.3. Installing Optional Packages | 22 |
| 1.5.4. Installing Debugging Information | 23 |
| 1.6. Updating Red Hat Developer Toolset | 23 |
| 1.6.1. Updating to a Minor Version | 23 |
| 1.6.2. Updating to a Major Version | 24 |
| 1.7. Uninstalling Red Hat Developer Toolset | 24 |
| 1.8. Additional Resources | 24 |
| Online Documentation | 24 |
| See Also | 25 |
| Part II. Integrated Development Environments | 26 |
| Chapter 2. Eclipse | 27 |
| 2.1. Installing Eclipse | 28 |
| 2.1.1. Building Eclipse from the Source RPM Packages | 29 |
| 2.2. Using Eclipse | 29 |
| 2.2.1. Using the Red Hat Developer Toolset Toolchain | 29 |
| 2.2.2. Using the Red Hat Enterprise Linux Toolchain | 30 |
| 2.3. Additional Resources | 31 |
| Installed Documentation | 31 |
| Online Documentation | 31 |
| See Also | 31 |
| Part III. Development Tools | 32 |
| Chapter 3. GNU Compiler Collection (GCC) | 33 |
| 3.1. GNU C Compiler | 33 |
| 3.1.1. Installing the C Compiler | 33 |
| 3.1.2. Using the C Compiler | 33 |
| 3.1.3. Running a C Program | 35 |
| 3.2. GNU C++ Compiler | 35 |
| 3.2.1. Installing the C++ Compiler | 35 |

| | |
|---|-----------|
| 3.2.2. Using the C++ Compiler | 35 |
| 3.2.3. Running a C++ Program | 37 |
| 3.3. GNU Fortran Compiler | 37 |
| 3.3.1. Installing the Fortran Compiler | 37 |
| 3.3.2. Using the Fortran Compiler | 37 |
| 3.3.3. Running a Fortran Program | 39 |
| 3.4. Additional Resources | 39 |
| Installed Documentation | 39 |
| Online Documentation | 40 |
| See Also | 40 |
| Chapter 4. binutils | 41 |
| 4.1. Installing binutils | 41 |
| 4.2. Using the GNU Assembler | 42 |
| 4.3. Using the GNU Linker | 42 |
| 4.4. Using Other Binary Tools | 43 |
| 4.5. Additional Resources | 43 |
| Installed Documentation | 44 |
| Online Documentation | 44 |
| See Also | 44 |
| Chapter 5. elfutils | 45 |
| 5.1. Installing elfutils | 45 |
| 5.2. Using elfutils | 45 |
| 5.3. Additional Resources | 46 |
| See Also | 46 |
| Chapter 6. dwz | 47 |
| 6.1. Installing dwz | 47 |
| 6.2. Using dwz | 47 |
| 6.3. Additional Resources | 47 |
| Installed Documentation | 48 |
| See Also | 48 |
| Part IV. Revision Control Systems | 49 |
| Chapter 7. Git | 50 |
| 7.1. Installing Git | 50 |
| 7.2. Configuring Git | 50 |
| 7.2.1. Configuring the Default Text Editor | 50 |
| 7.2.2. Setting Up User Information | 50 |
| 7.3. Using Git | 51 |
| 7.3.1. Creating a New Repository | 51 |
| Initializing an Empty Repository | 51 |
| Importing Data to a Repository | 51 |
| 7.3.2. Cloning an Existing Repository | 52 |
| Cloning a Git Repository | 52 |
| Cloning a Subversion Repository | 52 |
| Cloning a CVS Repository | 52 |
| 7.3.3. Adding, Renaming, and Deleting Files | 52 |
| Adding Files and Directories | 52 |
| Renaming Files and Directories | 53 |
| Deleting Files and Directories | 53 |
| 7.3.4. Viewing Changes | 53 |
| Viewing the Current Status | 53 |
| Viewing Differences | 53 |
| 7.3.5. Committing Changes | 54 |

| | |
|--|-----------|
| 7.3.6. Sharing Changes | 54 |
| Pushing Changes to a Public Git Repository | 54 |
| Creating Patches from Individual Commits | 54 |
| 7.3.7. Updating a Git Repository | 55 |
| 7.4. Additional Resources | 55 |
| Installed Documentation | 55 |
| Online Documentation | 55 |
| See Also | 55 |
| Part V. Debugging Tools | 57 |
| Chapter 8. GNU Debugger (GDB) | 58 |
| 8.1. Installing the GNU Debugger | 58 |
| 8.2. Preparing a Program for Debugging | 58 |
| Compiling Programs with Debugging Information | 58 |
| Installing Debugging Information for Existing Packages | 59 |
| 8.3. Running the GNU Debugger | 59 |
| 8.4. Listing Source Code | 60 |
| 8.5. Setting Breakpoints | 62 |
| Setting a New Breakpoint | 62 |
| Listing Breakpoints | 62 |
| Deleting Existing Breakpoints | 62 |
| 8.6. Starting Execution | 63 |
| 8.7. Displaying Current Values | 63 |
| 8.8. Continuing Execution | 64 |
| 8.9. Additional Resources | 65 |
| Online Documentation | 65 |
| See Also | 65 |
| Chapter 9. strace | 67 |
| 9.1. Installing strace | 67 |
| 9.2. Using strace | 67 |
| 9.2.1. Redirecting Output to a File | 68 |
| 9.2.2. Tracing Selected System Calls | 68 |
| 9.2.3. Displaying Time Stamps | 69 |
| 9.2.4. Displaying a Summary | 70 |
| 9.3. Additional Resources | 70 |
| Installed Documentation | 71 |
| See Also | 71 |
| Chapter 10. memstomp | 72 |
| 10.1. Installing memstomp | 74 |
| 10.2. Using memstomp | 74 |
| 10.3. Additional Resources | 76 |
| Installed Documentation | 76 |
| See Also | 76 |
| Part VI. Performance Monitoring Tools | 77 |
| Chapter 11. SystemTap | 78 |
| 11.1. Installing SystemTap | 78 |
| 11.2. Using SystemTap | 79 |
| 11.3. Additional Resources | 79 |
| Installed Documentation | 79 |
| Online Documentation | 80 |
| See Also | 80 |

| | |
|--|------------|
| Chapter 12. Valgrind | 81 |
| 12.1. Installing Valgrind | 81 |
| 12.2. Using Valgrind | 82 |
| 12.3. Additional Resources | 82 |
| Installed Documentation | 82 |
| Online Documentation | 83 |
| See Also | 83 |
| Chapter 13. OProfile | 84 |
| 13.1. Installing OProfile | 84 |
| 13.2. Using OProfile | 85 |
| 13.3. Additional Resources | 85 |
| Installed Documentation | 85 |
| Online Documentation | 86 |
| See Also | 86 |
| Chapter 14. Dyninst | 87 |
| 14.1. Installing Dyninst | 87 |
| 14.2. Using Dyninst | 87 |
| 14.2.1. Using Dyninst with SystemTap | 87 |
| 14.2.2. Using Dyninst as a Stand-alone Application | 89 |
| 14.3. Additional Resources | 94 |
| Installed Documentation | 94 |
| Online Documentation | 95 |
| See Also | 95 |
| Part VII. Getting Help | 96 |
| Chapter 15. Accessing Red Hat Product Documentation | 97 |
| Red Hat Developer Toolset | 97 |
| Red Hat Enterprise Linux | 97 |
| Chapter 16. Contacting Global Support Services | 98 |
| 16.1. Gathering Required Information | 98 |
| Background Information | 98 |
| Diagnostics | 98 |
| Account and Contact Information | 98 |
| Issue Severity | 99 |
| 16.2. Escalating an Issue | 99 |
| 16.3. Re-opening a Service Request | 100 |
| 16.4. Additional Resources | 100 |
| Online Documentation | 100 |
| Changes in Version 2.1 | 101 |
| A.1. Changes in Eclipse | 101 |
| A.1.1. Changes Since Red Hat Developer Toolset 2.0 | 101 |
| A.2. Changes in GCC | 101 |
| A.3. Changes in elfutils | 102 |
| A.3.1. Changes Since Red Hat Developer Toolset 2.0 | 102 |
| A.4. Changes in Git | 102 |
| A.4.1. Changes Since Red Hat Enterprise Linux 6.5 | 103 |
| General Changes in the Git Suite | 103 |
| The git am Command | 104 |
| The git apply Command | 104 |
| The git archive Command | 104 |
| The git bisect Command | 104 |
| The git blame Command | 105 |

| | |
|--|------------|
| The git branch Command | 105 |
| The git checkout Command | 105 |
| The git check-attr Command | 106 |
| The git cherry-pick Command | 106 |
| The git clean Command | 106 |
| The git clone Command | 106 |
| The git commit Command | 107 |
| The git count-objects Command | 107 |
| The git cvsimport Command | 107 |
| The git cvsserver Command | 108 |
| The git daemon Command | 108 |
| The git describe Command | 108 |
| The git diff Command | 108 |
| The git fetch Command | 109 |
| The git format-patch Command | 109 |
| The git for-each-ref Command | 110 |
| The git grep Command | 110 |
| The git hash-object Command | 111 |
| The git help Command | 111 |
| The git imap-send Command | 111 |
| The git init Command | 111 |
| The git log Command | 111 |
| The git ls-remote Command | 112 |
| The git merge Command | 112 |
| The git mergetool Command | 113 |
| The git merge-base Command | 113 |
| The git merge-file Command | 113 |
| The git notes Command | 113 |
| The git notes prune Subcommand | 113 |
| The git pull Command | 114 |
| The git push Command | 114 |
| The git read-tree Command | 114 |
| The git rebase Command | 114 |
| The git remote Command | 115 |
| The git replace Command | 115 |
| The git rerere Command | 115 |
| The git reset Command | 115 |
| The git revert Command | 115 |
| The git send-email Command | 115 |
| The git show Command | 116 |
| The git show-ref Command | 116 |
| The git stash Command | 116 |
| The git status Command | 116 |
| The git submodule Command | 117 |
| The git tag Command | 117 |
| The gitweb Interface | 117 |
| Other new subcommands and exotic option combinations | 117 |
| A.5. Changes in GDB | 118 |
| A.5.1. Changes Since Red Hat Developer Toolset 2.0 | 118 |
| A.6. Changes in Valgrind | 118 |
| A.6.1. Changes Since Red Hat Developer Toolset 2.0 | 118 |
| A.7. Changes in OProfile | 118 |
| A.7.1. Changes Since Red Hat Developer Toolset 2.0 | 118 |
| Changes in Version 2.0 | 119 |
| B.1. Changes in Eclipse | 119 |

| | |
|---|-----|
| B.1.1. Changes Since Red Hat Enterprise Linux 6.4 | 119 |
| B.2. Changes in GCC | 120 |
| B.2.1. Changes Since Red Hat Developer Toolset 1.1 | 120 |
| B.2.1.1. Caveats | 120 |
| Aggressive Loop Optimizations | 120 |
| B.2.1.2. General Improvements and Changes | 121 |
| New Local Register Allocator | 121 |
| AddressSanitizer | 121 |
| ThreadSanitizer | 121 |
| Compiling Extremely Large Functions | 121 |
| New -Og Optimization Level | 121 |
| Caret Diagnostic Messages | 121 |
| New -fira-hoist-pressure Option | 122 |
| New -fopt-info Option | 122 |
| New -floop-nest-optimize Option | 122 |
| Hot and Cold Attributes on Labels | 122 |
| B.2.1.3. Debugging Enhancements | 122 |
| DWARF4 | 122 |
| New -gsplit-dwarf Option | 122 |
| B.2.1.4. C++ Changes | 123 |
| Experimental C++ Features from an Upcoming Standard | 123 |
| New thread_local Keyword | 123 |
| Dynamic Initialization of Thread-local Variables | 123 |
| C++11 Attribute Syntax | 123 |
| C++11 Alignment Specifier | 124 |
| B.2.1.5. Fortran Changes | 124 |
| B.2.1.5.1. Caveats | 124 |
| B.2.1.5.2. ABI Compatibility | 124 |
| B.2.1.5.3. Other Changes | 124 |
| BACKTRACE Intrinsic | 124 |
| Floating Point Numbers with “q” as Exponential | 124 |
| GFORTTRAN_TMPDIR Environment Variable | 124 |
| Fortran 2003 | 125 |
| TS 29113 | 125 |
| B.2.1.6. x86-specific Improvements | 125 |
| New Instructions | 125 |
| New Built-in Functions to Detect Run-time CPU Type and ISA | 125 |
| Function Multiversioning | 126 |
| New RTM and HLE Intrinsics | 126 |
| Transactions Using Transactional Synchronization Extensions | 127 |
| Support for AMD Family 15h Processors | 127 |
| Support for AMD Family 16h Processors | 127 |
| B.2.2. Changes Since Red Hat Enterprise Linux 6.4 and 5.9 | 127 |
| B.2.2.1. Status and Features | 127 |
| B.2.2.1.1. C++11 | 127 |
| B.2.2.1.2. C11 | 128 |
| B.2.2.1.3. Parallelism and Concurrency | 128 |
| C++11 Types and GCC Built-ins for Atomic Memory Access | 128 |
| Transactional Memory | 129 |
| B.2.2.1.4. Architecture-specific Options | 130 |
| B.2.2.1.5. Link-time Optimization | 132 |
| B.2.2.1.6. Miscellaneous | 132 |
| B.2.2.2. Language Compatibility | 133 |
| B.2.2.2.1. C | 133 |
| Duplicate Member | 133 |

| | |
|--|-----|
| B.2.2.2.2. C++ | 133 |
| Header Dependency Changes | 133 |
| Name Lookup Changes | 134 |
| Uninitialized const | 135 |
| Visibility of Template Instantiations | 135 |
| User-defined Literal Support | 135 |
| Taking the Address of Temporary | 136 |
| Miscellaneous | 136 |
| B.2.2.2.3. C/C++ Warnings | 136 |
| B.2.2.2.4. Fortran | 137 |
| B.2.2.2.4.1. New Features | 137 |
| B.2.2.2.4.2. Compatibility Changes | 138 |
| B.2.2.2.4.3. Fortran 2003 Features | 139 |
| B.2.2.2.4.4. Fortran 2003 Compatibility | 139 |
| B.2.2.2.4.5. Fortran 2008 Features | 140 |
| B.2.2.2.4.6. Fortran 2008 Compatibility | 141 |
| B.2.2.2.4.7. Fortran 77 Compatibility | 141 |
| B.2.2.3. ABI Compatibility | 141 |
| B.2.2.3.1. C++ ABI | 141 |
| B.2.2.3.2. Miscellaneous | 142 |
| B.2.2.4. Debugging Compatibility | 142 |
| B.2.2.5. Other Compatibility | 142 |
| B.3. Changes in binutils | 142 |
| B.3.1. GNU Linker | 143 |
| B.3.1.1. New Features | 143 |
| Changes Since Red Hat Enterprise Linux 6.4 | 143 |
| Changes Since Red Hat Enterprise Linux 5.9 | 143 |
| B.3.1.2. Compatibility Changes | 144 |
| Changes Since Red Hat Enterprise Linux 6.4 | 144 |
| Changes Since Red Hat Enterprise Linux 5.9 | 144 |
| B.3.2. GNU Assembler | 145 |
| B.3.2.1. New Features | 145 |
| Changes Since Red Hat Enterprise Linux 6.4 | 145 |
| Changes Since Red Hat Enterprise Linux 5.9 | 145 |
| B.3.3. Other Binary Tools | 146 |
| B.3.3.1. New Features | 146 |
| Changes Since Red Hat Developer Toolset 1.1 | 146 |
| Changes Since Red Hat Enterprise Linux 6.4 | 146 |
| Changes Since Red Hat Enterprise Linux 5.9 | 146 |
| B.3.3.2. Compatibility Changes | 147 |
| Changes Since Red Hat Enterprise Linux 5.9 | 147 |
| B.4. Changes in elfutils | 147 |
| B.4.1. Changes Since Red Hat Developer Toolset 1.1 | 147 |
| B.4.2. Changes Since Red Hat Enterprise Linux 6.4 | 148 |
| B.4.3. Changes Since Red Hat Enterprise Linux 5.9 | 148 |
| B.5. Changes in dwz | 149 |
| B.5.1. Changes Since Red Hat Developer Toolset 1.1 | 149 |
| B.6. Changes in GDB | 149 |
| B.6.1. Changes Since Red Hat Developer Toolset 1.1 | 149 |
| B.6.2. Changes Since Red Hat Enterprise Linux 6.4 | 152 |
| New Features | 152 |
| Compatibility Changes | 156 |
| B.6.3. Changes Since Red Hat Enterprise Linux 5.9 | 157 |
| New Features | 157 |
| B.7. Changes in strace | 160 |

| | |
|---|------------|
| B.7.1. Changes Since Red Hat Enterprise Linux 6.4 and 5.9 | 160 |
| B.8. Changes in SystemTap | 160 |
| B.8.1. Changes Since Red Hat Developer Toolset 1.1 | 160 |
| B.9. Changes in OProfile | 161 |
| B.9.1. Changes Since Red Hat Developer Toolset 1.1 | 161 |
| B.9.2. Changes Since Red Hat Enterprise Linux 5.9 | 161 |
| B.10. Changes in Valgrind | 162 |
| B.10.1. Changes Since Red Hat Developer Toolset 1.1 | 162 |
| B.10.2. Changes Since Red Hat Enterprise Linux 5.9 | 162 |
| Revision History | 164 |
| Index | 164 |
| A | 164 |
| B | 164 |
| C | 165 |
| D | 165 |
| E | 166 |
| F | 168 |
| G | 168 |
| H | 170 |
| L | 170 |
| M | 170 |
| N | 170 |
| O | 170 |
| R | 172 |
| S | 172 |
| V | 174 |

Preface

1. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

In PDF and paper editions, this manual uses typefaces drawn from the [Liberation Fonts](#) set. The Liberation Fonts set is also used in HTML editions if the set is installed on your system. If not, alternative but equivalent typefaces are displayed. Note: Red Hat Enterprise Linux 5 and later include the Liberation Fonts set by default.

1.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

Mono-spaced Bold

Used to highlight system input, including shell commands, file names and paths. Also used to highlight keys and key combinations. For example:

To see the contents of the file **my_next_bestselling_novel** in your current working directory, enter the **cat my_next_bestselling_novel** command at the shell prompt and press **Enter** to execute the command.

The above includes a file name, a shell command and a key, all presented in mono-spaced bold and all distinguishable thanks to context.

Key combinations can be distinguished from an individual key by the plus sign that connects each part of a key combination. For example:

Press **Enter** to execute the command.

Press **Ctrl+Alt+F2** to switch to a virtual terminal.

The first example highlights a particular key to press. The second example highlights a key combination: a set of three keys pressed simultaneously.

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in **mono-spaced bold**. For example:

File-related classes include **filesystem** for file systems, **file** for files, and **dir** for directories. Each class has its own associated set of permissions.

Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialog-box text; labeled buttons; check-box and radio-button labels; menu titles and submenu titles. For example:

Choose **System** → **Preferences** → **Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, select the **Left-handed mouse** check box and click **Close** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).

To insert a special character into a **gedit** file, choose **Applications** → **Accessories** →

Character Map from the main menu bar. Next, choose **Search** → **Find...** from the **Character Map** menu bar, type the name of the character in the **Search** field and click **Next**. The character you sought will be highlighted in the **Character Table**. Double-click this highlighted character to place it in the **Text to copy** field and then click the **Copy** button. Now switch back to your document and choose **Edit** → **Paste** from the **gedit** menu bar.

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in proportional bold and all distinguishable by context.

Mono-spaced Bold Italic** or **Proportional Bold Italic

Whether mono-spaced bold or proportional bold, the addition of italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

To connect to a remote machine using ssh, type **ssh *username@domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh *john@example.com***.

The **mount -o remount *file-system*** command remounts the named file system. For example, to remount the **/home** file system, the command is **mount -o remount */home***.

To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: ***package-version-release***.

Note the words in bold italics above: *username*, *domain.name*, *file-system*, *package*, *version* and *release*. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, italics denotes the first use of a new and important term. For example:

Publican is a *DocBook* publishing system.

1.2. Pull-quote Conventions

Terminal output and source code listings are set off visually from the surrounding text.

Output sent to a terminal is set in **mono-spaced roman** and presented thus:

```
books      Desktop  documentation  drafts  mss    photos  stuff  svn
books_tests Desktop1  downloads      images  notes  scripts svgs
```

Source-code listings are also set in **mono-spaced roman** but add syntax highlighting as follows:

```
static int kvm_vm_ioctl_deassign_device(struct kvm *kvm,
                                       struct kvm_assigned_pci_dev *assigned_dev)
{
    int r = 0;
    struct kvm_assigned_dev_kernel *match;

    mutex_lock(&kvm->lock);

    match = kvm_find_assigned_dev(&kvm->arch.assigned_dev_head,
                                assigned_dev->assigned_dev_id);
    if (!match) {
        printk(KERN_INFO "%s: device hasn't been assigned before, "
                    "so cannot be deassigned\n", __func__);
        r = -EINVAL;
        goto out;
    }

    kvm_deassign_device(kvm, match);

    kvm_free_assigned_device(kvm, match);

out:
    mutex_unlock(&kvm->lock);
    return r;
}
```

1.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.



Note

Notes are tips, shortcuts or alternative approaches to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.



Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled “Important” will not cause data loss but may cause irritation and frustration.



Warning

Warnings should not be ignored. Ignoring warnings will most likely cause data loss.

2. Getting Help and Giving Feedback

2.1. Do You Need Help?

If you experience difficulty with a procedure described in this documentation, visit the Red Hat Customer

Portal at <http://access.redhat.com>. Through the customer portal, you can:

- search or browse through a knowledgebase of technical support articles about Red Hat products.
- submit a support case to Red Hat Global Support Services (GSS).
- access other product documentation.

Red Hat also hosts a large number of electronic mailing lists for discussion of Red Hat software and technology. You can find a list of publicly available mailing lists at <https://www.redhat.com/mailman/listinfo>. Click on the name of any mailing list to subscribe to that list or to access the list archives.

2.2. We Need Feedback!

If you find a typographical error in this manual, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report in Bugzilla: <http://bugzilla.redhat.com/> against the product **Red Hat Developer Toolset**.

When submitting a bug report, be sure to mention the manual's identifier: *doc-User_Guide*

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, please include the section number and some of the surrounding text so we can find it easily.

Part I. Introduction

Chapter 1. Red Hat Developer Toolset

1.1. About Red Hat Developer Toolset

Red Hat Developer Toolset is a Red Hat offering for developers on the Red Hat Enterprise Linux platform, and provides a complete set of development and performance analysis tools that can be installed and used on multiple versions of Red Hat Enterprise Linux. Executables built with the Red Hat Developer Toolset toolchain can then also be deployed and run on multiple versions of Red Hat Enterprise Linux. For detailed compatibility information, see [Section 1.3, “Compatibility”](#).

Red Hat Developer Toolset does not replace the default system tools provided with Red Hat Enterprise Linux 5 or 6 when installed on those platforms. Instead, a parallel set of developer tools provides an alternative, newer version of those tools for optional use by developers. The default compiler and debugger, for example, remain those provided by the base Red Hat Enterprise Linux system.

Table 1.1. Red Hat Developer Toolset Components

| Name | Version | Platforms | Description |
|-----------|---------|---|---|
| Eclipse | 4.3.1 | Red Hat Enterprise Linux 6 | An integrated development environment for the graphical user interface. ^[a] |
| GCC | 4.8.2 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A portable compiler suite with support for C, C++, and Fortran. |
| binutils | 2.23.52 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A collection of binary tools and other utilities to inspect and manipulate object files and binaries. |
| elfutils | 0.157 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A collection of binary tools and other utilities to inspect and manipulate ELF files. |
| dwz | 0.11 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A tool to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size. |
| Git | 1.8.4 | Red Hat Enterprise Linux 6 | A distributed revision control system with a peer-to-peer architecture. |
| GDB | 7.6.1 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A command line debugger for programs written in C, C++, and Fortran. |
| strace | 4.7 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A debugging tool to monitor system calls that a program uses and signals it receives. |
| memstomp | 0.1.4 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A debugging tool to identify calls to library functions with overlapping memory regions that are not allowed by various standards. |
| SystemTap | 2.1 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A tracing and probing tool to monitor the activities of the entire system without the need to instrument, recompile, install, and reboot. |
| Valgrind | 3.8.1 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | An instrumentation framework and a number of tools to profile applications in order to detect memory errors, identify memory management problems, and report any use of improper arguments in system calls. |
| OProfile | 0.9.8 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A system-wide profiler that uses the performance monitoring hardware on the processor to retrieve information about the kernel and executables on the system. |
| Dyninst | 8.0 | Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6 | A library for instrumenting and working with user-space executables during their execution. |

[a] If you intend to develop applications for Red Hat JBoss Middleware or require support for OpenShift Tools, it is recommended that you use [Red Hat JBoss Developer Studio](#).

Red Hat Developer Toolset differs from “[Technology Preview](#)” compiler releases previously supplied in

Red Hat Enterprise Linux in two important respects:

1. Red Hat Developer Toolset can be used on multiple major and minor releases of Red Hat Enterprise Linux, as detailed in [Section 1.3, “Compatibility”](#).
2. Unlike Technology Preview compilers and other tools shipped in earlier Red Hat Enterprise Linux, Red Hat Developer Toolset is fully supported under Red Hat Enterprise Linux Subscription Level Agreements, is functionally complete, and is intended for production use.

Important bug fixes and security errata are issued to Red Hat Developer Toolset subscribers in a similar manner to Red Hat Enterprise Linux for two years from the release of each major version release. A new major version of Red Hat Developer Toolset is released annually, providing significant updates for existing components and adding major new components. A single minor release, issued six months after each new major version release, provides a smaller update of bug fixes, security errata, and new minor components.

Additionally, the Red Hat Enterprise Linux Application Compatibility Specification also applies to Red Hat Developer Toolset (subject to some constraints on the use of newer C++11 language features, detailed in [Section B.2.2.3, “ABI Compatibility”](#)).



Important

Applications and libraries provided by Red Hat Developer Toolset do not replace the Red Hat Enterprise Linux system versions, nor are they used in preference to the system versions. Using a framework called **Software Collections**, an additional set of developer tools is installed into the `/opt` directory and is explicitly enabled by the user on demand using the supplied `sc1` utility.

1.2. Main Features

The Red Hat Developer Toolset version of the **GNU Compiler Collection (GCC)** provides the following features:

- A new register allocator (LRA) has been added, improving code performance.
- A fast memory error detector called AddressSanitizer has been added.
- A fast data race detector called ThreadSanitizer has been added.
- Extremely large functions can now be compiled faster using less memory.
- A new **general** optimization level has been introduced.
- GCC diagnostic messages now highlight the exact problem in the source code.
- Various new optimization options have been added.
- DWARF4 is now used as the default debug format.
- GCC now fully implements the C++11 language standard.
- C++11 library support has been extended though is still experimental.
- GCC now supports dynamic initialization of thread-local variables.
- Support has been added for Intel **FXSR**, **XSAVE**, and **XSAVEOPT** instructions.
- New built-in functions added to detect run-time Intel CPU Type and ISA.
- Intel function multi-versioning support added.
- Intel **RTM/HLE** intrinsics, built-ins, and code generation have been added.
- Transactions (the **-fgnu-tm** option) can now be run using Intel **TSX** extensions.
- Support for AMD family 15h and 16h processors has been added.

- Various Fortran changes have been included.

The version of the **GNU Debugger (GDB)** included in Red Hat Developer Toolset provides the following features:

- Improved and expanded support for Python scripting.
- Improved handling of C++ debuggee executables.
- Improved inferior control commands.
- Improved support for ambiguous line specifications.
- Improved tracepoint support.
- Multi-program debugging.

Additionally, the Red Hat Developer Toolset version of **binutils** provides these features:

- The new **gold** linker, which is smaller and faster than **ld**. Note that **gold** is not the default linker and must be explicitly enabled by using the **alternatives** command.
- Support for link-time optimization (LTO) in conjunction with GCC.
- Support for build-IDs, unique numbers to identify executables.
- Support for the **IFUNC** and **UNIQUE** symbols that are used by **glibc** to improve performance. Due to dependencies on a particular version of the **glibc** library, these symbols are only available on Red Hat Enterprise Linux 6.
- Compressed debug sections for smaller debug info files.

For a full list of changes and features introduced in this release, see [Appendix A, Changes in Version 2.1](#).

1.3. Compatibility

Red Hat Developer Toolset 2.1 is available for Red Hat Enterprise Linux 5 and 6, both for 32-bit and 64-bit Intel and AMD architectures. [Figure 1.1, “Red Hat Developer Toolset 2.1 Compatibility Matrix”](#) illustrates the support for binaries built with Red Hat Developer Toolset on a certain version of Red Hat Enterprise Linux when those binaries are run on various other versions of this system.

For ABI compatibility information, see [Section B.2.2.3, “ABI Compatibility”](#).

| | | “Run on” | | | | |
|---------------|---------|----------|------|---------|---------|-----|
| | | 5.9 EUS | 5.10 | 6.3 EUS | 6.4 EUS | 6.5 |
| “Built with” | 5.9 EUS | ✓ | ✓ | ✓ | ✓ | ✓ |
| | 5.10 | ✗ | ✓ | ✓ | ✓ | ✓ |
| | 6.3 EUS | ✗ | ✗ | ✓ | ✓ | ✓ |
| | 6.4 EUS | ✗ | ✗ | ✗ | ✓ | ✓ |
| | 6.5 | ✗ | ✗ | ✗ | ✗ | ✓ |
| ✗ Unsupported | | | | | | |
| ✓ Supported | | | | | | |

Figure 1.1. Red Hat Developer Toolset 2.1 Compatibility Matrix

1.4. Getting Access to Red Hat Developer Toolset

Red Hat Developer Toolset is an offering that is distributed as part of the developer subscriptions listed in [Table 1.2, “Subscriptions That Provide Red Hat Developer Toolset”](#). Depending on the subscription management service with which you registered your Red Hat Enterprise Linux system, you can either enable Red Hat Developer Toolset by using the Red Hat Subscription Management, or by using RHN Classic.

For detailed instructions on how to enable Red Hat Developer Toolset using RHN Classic or Red Hat Subscription Management, see the respective section below. For information on how to register your system with one of these subscription management services, see the *Red Hat Subscription Management Guide*.



Important

If you are running a version of Red Hat Enterprise Linux prior to 5.9 or 6.4, you will be unable to download Red Hat Developer Toolset through Red Hat Subscription Management. To obtain Red Hat Developer Toolset, you will need to either update to Red Hat Enterprise Linux 5.9 or 6.4, or register your system with RHN Classic. For more information, see <https://access.redhat.com/site/solutions/129003>.

Table 1.2. Subscriptions That Provide Red Hat Developer Toolset

| SKU | Subscription Name |
|-----------|---|
| RH2282403 | Red Hat Enterprise Linux Developer Support, Professional |
| RH2264664 | Red Hat Enterprise Linux Developer Support, Enterprise |
| RH2262474 | Red Hat Enterprise Linux Developer Suite |
| RH3482727 | Red Hat Enterprise Linux Developer Workstation, Professional |
| RH3413336 | Red Hat Enterprise Linux Developer Workstation, Enterprise |
| RH3414706 | 30 day Self-Supported Red Hat Enterprise Linux Developer Workstation Evaluation |
| RH3474212 | 60 day Supported Red Hat Enterprise Linux Developer Workstation Evaluation |
| RH3437268 | 90 day Supported Red Hat Enterprise Linux Developer Workstation Evaluation |
| SER0402 | 1-year Unsupported Partner Evaluation Red Hat Enterprise Linux |
| SER0403 | 1-year Unsupported Red Hat Advanced Partner Subscription |

1.4.1. Using RHN Classic

If your system is registered with RHN Classic, complete the following steps to subscribe to Red Hat Developer Toolset:

1. Determine the exact name of the Red Hat Developer Toolset channel. To do so, type the following at a shell prompt as **root** to display a list of all channels that are available to you:

```
rhn-channel --available-channels
```

The name of the channel depends on the specific version of Red Hat Enterprise Linux you are using and is in the following format:

```
rhel-architecture-variant-dts2-version
```

In addition, certain packages such as *devtoolset-2-git* or *devtoolset-2-gcc-plugin-devel* depend on packages that are only available in the **Optional** channel. The name of this channel uses the following format:

```
rhel-architecture-variant-optional-version
```

Replace ***architecture*** with the system's CPU architecture (**x86_64** or **i386**), ***variant*** with the Red Hat Enterprise Linux system variant (**server** or **workstation**), and ***version*** with the Red Hat Enterprise Linux system version (**5** or **6**).

2. Subscribe the system to the channels from [Step 1](#) by running the following command as **root**:

```
rhnc-channel --add --channel=channel_name
```

Replace ***channel_name*** with the name of the channel to enable.

3. To verify the list of channels you are subscribed to, at any time, run as **root**:

```
rhnc-channel --list
```

Once the system is subscribed, you can install Red Hat Developer Toolset as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#). For more information on how to register your system with RHN Classic, see the *Red Hat Subscription Management Guide*.

1.4.2. Using Red Hat Subscription Management

If your system is registered with Red Hat Subscription Management, complete the following steps to attach a subscription that provides access to the repository for Red Hat Developer Toolset, and then enable that repository:

1. Determine the pool ID of a subscription that provides Red Hat Developer Toolset. To do so, type the following at a shell prompt as **root** to display a list of all subscriptions that are available for your system:

```
subscription-manager list --available
```

For each available subscription, this command displays its name, unique identifier, expiration date, and other details related to your subscription. The pool ID is listed on a line beginning with **Pool Id**.

For a complete list of subscriptions that provide access to Red Hat Developer Toolset, see [Table 1.2, “Subscriptions That Provide Red Hat Developer Toolset”](#).

2. Attach the appropriate subscription to your system by running the following command as **root**:

```
subscription-manager subscribe --pool=pool_id
```

Replace ***pool_id*** with the pool ID you determined in the previous step. To verify the list of subscriptions your system has currently attached, at any time, run as **root**:

```
subscription-manager list --consumed
```

- Determine the exact name of Red Hat Developer Toolset repositories. To do so, type the following at a shell prompt as **root** to retrieve repository metadata and to display a list of available Yum repositories:

```
subscription-manager repos --list
```

The repository names depend on the specific version of Red Hat Enterprise Linux you are using, and are in the following format:

```
rhel-variant-dts2-version-rpms
rhel-variant-dts2-version-debug-rpms
rhel-variant-dts2-version-source-rpms
```

In addition, certain packages such as *devtoolset-2-git* or *devtoolset-2-gcc-plugin-devel* depend on packages that are only available in the **Optional** channel. The repository names with these packages use the following format:

```
rhel-version-variant-optional-rpms
rhel-version-variant-optional-debug-rpms
rhel-version-variant-optional-source-rpms
```

Replace **variant** with the Red Hat Enterprise Linux system variant (**server** or **workstation**), and **version** with the Red Hat Enterprise Linux system version (**5** or **6**).

- Enable the repositories from [Step 3](#) by running the following command as **root**:

```
subscription-manager repos --enable repository
```

Replace **repository** with the name of the repository to enable.

Once the subscription is attached to the system, you can install Red Hat Developer Toolset as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#). For more information on how to register your system using Red Hat Subscription Management and associate it with subscriptions, see the *Red Hat Subscription Management Guide*.

1.5. Installing Red Hat Developer Toolset

Red Hat Developer Toolset is distributed as a collection of RPM packages that can be installed, updated, uninstalled, and inspected by using the standard package management tools that are included in Red Hat Enterprise Linux. Note that a valid subscription is required in order to install Red Hat Developer Toolset on your system. For detailed instructions on how to associate your system with an appropriate subscription and get access to the product, see [Section 1.4, “Getting Access to Red Hat Developer Toolset”](#).



Important

Before installing Red Hat Developer Toolset on a system that is already running the previous version of the product, make sure that the *devtoolset-1.1-gcc-debuginfo* is not installed. To uninstall this package from the system, type the following at a shell prompt as **root**:

```
yum remove devtoolset-1.1-gcc-debuginfo
```

**Important**

After installing Red Hat Developer Toolset, it is recommended to apply all available Red Hat Enterprise Linux errata updates to enable all Red Hat Developer Toolset features and apply fixes that may otherwise impact the tools or built code.

1.5.1. Installing All Available Components

To install all components that are included in this product, install the *devtoolset-2* package by typing the following at a shell prompt as **root**:

```
yum install devtoolset-2
```

This installs the Eclipse development environment, all development, debugging, and performance monitoring tools, the **scl** utility and other dependent packages to the system. Alternatively, you can choose to install only a selected package group as described in [Section 1.5.2, “Installing Individual Package Groups”](#).

1.5.2. Installing Individual Package Groups

To make it easier to install only certain components such as the integrated development environment or the software development toolchain, Red Hat Developer Toolset is distributed with a number of meta packages that allow you to install selected package groups as described in [Table 1.3, “Red Hat Developer Toolset Meta Packages”](#).

Table 1.3. Red Hat Developer Toolset Meta Packages

| Package Name | Description | Installed Components |
|-------------------------------|------------------------------------|---|
| <i>devtoolset-2-ide</i> | Integrated Development Environment | Eclipse |
| <i>devtoolset-2-perftools</i> | Performance monitoring tools | SystemTap, Valgrind, OProfile, Dyninst |
| <i>devtoolset-2-toolchain</i> | Development and debugging tools | GCC, GDB, binutils, elfutils, dwz, memstomp, strace |
| <i>devtoolset-2-vc</i> | Revision control systems | Git |

To install any of these meta packages, type the following at a shell prompt as **root**:

```
yum install package_name...
```

Replace **package_name** with a space-separated list of meta packages you want to install. For example, to install only the Eclipse development environment and packages that depend on it, type as **root**:

```
~]# yum install devtoolset-2-ide
```

Alternatively, you can choose to install all available components as described in [Section 1.5.1, “Installing All Available Components”](#).

1.5.3. Installing Optional Packages

Red Hat Developer Toolset is distributed with a number of optional packages that are not installed by

default. To list all Red Hat Developer Toolset packages that are available to you but not installed on your system, type the following command at a shell prompt:

```
yum list available devtoolset-2-\*
```

To install any of these optional packages, run as **root**:

```
yum install package_name...
```

Replace *package_name* with a space-separated list of packages that you want to install. For example, to install the *devtoolset-2-gdb-gdbserver* and *devtoolset-2-gdb-doc* packages, type:

```
~]# yum install devtoolset-2-gdb-gdbserver devtoolset-2-gdb-doc
```

1.5.4. Installing Debugging Information

To install debugging information for any of the Red Hat Developer Toolset packages, make sure that the *yum-utils* package is installed and run the following command as **root**:

```
debuginfo-install package_name
```

For example, to install debugging information for the *devtoolset-2-dwz* package, type:

```
~]# debuginfo-install devtoolset-2-dwz
```

Note that in order to use this command, you need to have access to the repository with these packages. If your system is registered with Red Hat Subscription Management, enable the **rhel-variant-dts2-version-debug-rpms** repository as described in [Section 1.4.2, “Using Red Hat Subscription Management”](#). If your system is registered with RHN Classic, subscribe the system to the **rhel-architecture-variant-version-debuginfo** channel as described in [Section 1.4.1, “Using RHN Classic”](#). For more information on how to get access to debuginfo packages, see <https://access.redhat.com/site/solutions/9907>.

1.6. Updating Red Hat Developer Toolset

1.6.1. Updating to a Minor Version

When a new minor version of Red Hat Developer Toolset is available, run the following command as **root** to update your Red Hat Enterprise Linux installation:

```
yum update
```

This updates all packages on your Red Hat Enterprise Linux system, including the Red Hat Developer Toolset versions the Eclipse development environment, development, debugging, and performance monitoring tools, the **sc1** utility and other dependent packages.



Important

Use of Red Hat Developer Toolset requires the removal of any earlier pre-release versions of this product. Additionally, it is not possible to update to Red Hat Developer Toolset 2.1 from a pre-release version of Red Hat Developer Toolset, including beta releases. If you have previously installed any pre-release version of Red Hat Developer Toolset, uninstall it from your system as described in [Section 1.7, “Uninstalling Red Hat Developer Toolset”](#) and install the new version as documented in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

1.6.2. Updating to a Major Version

When a new major version of Red Hat Developer Toolset is available, you can install it in parallel with the previous version of the product. For detailed instructions on how to install Red Hat Developer Toolset on your system, see [Section 1.5, “Installing Red Hat Developer Toolset”](#).

1.7. Uninstalling Red Hat Developer Toolset

To uninstall Red Hat Developer Toolset packages from your system, type the following at a shell prompt as **root**:

```
yum remove devtoolset-2\*
```

This removes the GNU Compiler Collection, GNU Debugger, binutils, and other packages that are part of Red Hat Developer Toolset from the system. To uninstall the **scl** utility as well, type as **root**:

```
yum remove scl-utils\*
```

Note that uninstallation of the tools provided by Red Hat Developer Toolset does not affect the Red Hat Enterprise Linux system versions of these tools.

For information on how to uninstall Red Hat Developer Toolset 1.1, see the [Red Hat Developer Toolset 1.1 User Guide](#).

1.8. Additional Resources

For more information about Red Hat Developer Toolset and Red Hat Enterprise Linux, see the resources listed below.

Online Documentation

- [Red Hat Subscription Management Guide](#) — The *Red Hat Subscription Management Guide* provides detailed information on how to manage subscriptions on Red Hat Enterprise Linux.
- [Red Hat Developer Toolset 2.1 Release Notes](#) — The *Release Notes* for Red Hat Developer Toolset 2.1 contain more information about this product.
- [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information on the Eclipse IDE, libraries and runtime support, compiling and building, debugging, and profiling on this system.
- [Red Hat Enterprise Linux 6 Installation Guide](#) — The *Installation Guide* for Red Hat Enterprise Linux 6 explains how to obtain, install, and update the system.
- [Red Hat Enterprise Linux 5 Installation Guide](#) — The *Installation Guide* for Red Hat Enterprise Linux

5 explains how to obtain, install, and update the system.

- [Red Hat Enterprise Linux 6 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 6 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 6.
- [Red Hat Enterprise Linux 5 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 5 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 5.

See Also

- [Appendix A, *Changes in Version 2.1*](#) provides a comprehensive list of changes and improvements over the Red Hat Enterprise Linux system versions of the GNU Compiler Collection, GNU Debugger, and binutils, as well as information about the language, ABI, and debugging compatibility.

Part II. Integrated Development Environments

Chapter 2. Eclipse

Eclipse is a powerful development environment that provides tools for each phase of the development process. It integrates a variety of disparate tools into a unified environment to create a rich development experience, provides a fully configurable user interface, and features a pluggable architecture that allows for extension in a variety of ways. For instance, the Valgrind plug-in allows programmers to perform memory profiling, otherwise performed on the command line, through the Eclipse user interface.

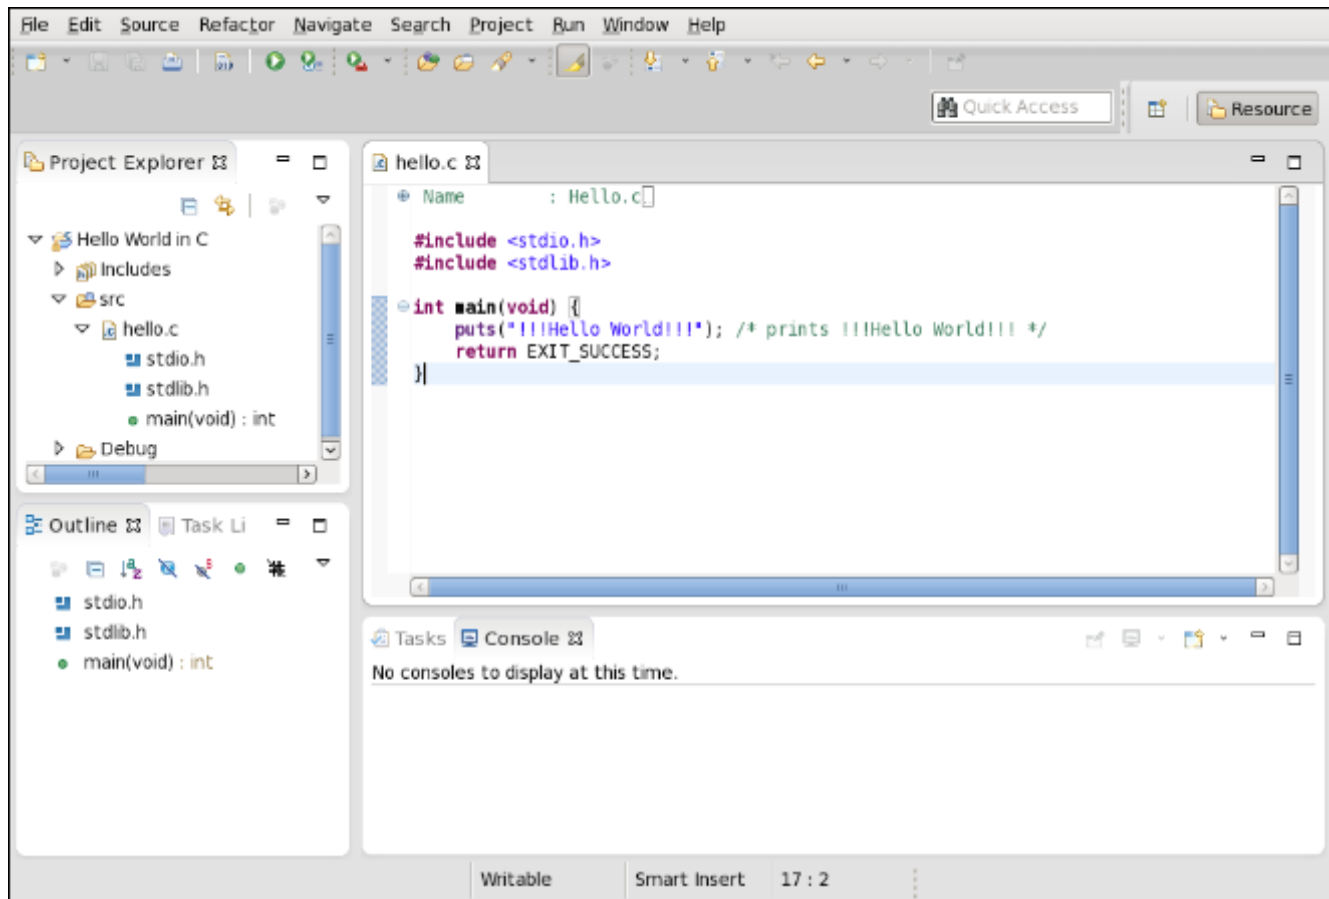


Figure 2.1. Sample Eclipse Session

Eclipse provides a graphical development environment alternative to traditional interaction with command line tools and as such, it is a welcome alternative to developers who do not want to use the command line interface. The traditional, mostly command line based Linux tools suite (such as **gcc** or **gdb**) and Eclipse offer two distinct approaches to programming.

Red Hat Developer Toolset is distributed with **Eclipse 4.3.1**, which is based on the Eclipse Foundation's 2013 Eclipse Kepler 4.3 Service Release 1. Note that if you intend to develop applications for Red Hat JBoss Middleware or require support for OpenShift Tools, it is recommended that you use [Red Hat JBoss Developer Studio](#).

Table 2.1. Eclipse Components Included in Red Hat Developer Toolset

| Package | Description |
|---|---|
| <i>devtoolset-2-eclipse-cdt</i> | The C/C++ Development Tooling (CDT), which provides features and plug-ins for development in C and C++. |
| <i>devtoolset-2-eclipse-emf</i> | The Eclipse Modeling Framework (EMF), which allows you to build applications based on a structured data model. |
| <i>devtoolset-2-eclipse-gef</i> | The Graphical Editing Framework (GEF), which allows you to create a rich graphical editor from an existing application model. |
| <i>devtoolset-2-eclipse-rse</i> | The Remote System Explorer (RSE) framework, which allows you to work with remote systems from Eclipse. |
| <i>devtoolset-2-eclipse-jgit</i> | JGit, a Java implementation of the Git revision control system. |
| <i>devtoolset-2-eclipse-egit</i> | EGit, a team provider for Eclipse that provides features and plug-ins for interaction with Git repositories. |
| <i>devtoolset-2-eclipse-myllyn</i> | Myllyn, a task management system for Eclipse. |
| <i>devtoolset-2-eclipse-linuxtools</i> | A meta package for Linux-specific Eclipse plug-ins. |
| <i>devtoolset-2-eclipse-changelog</i> ^[a] | The ChangeLog plug-ins, which allows you to create and maintain changelog files. |
| <i>devtoolset-2-eclipse-gcov</i> ^[a] | The GCov plug-in, which integrates the GCov test coverage program with Eclipse. |
| <i>devtoolset-2-eclipse-gprof</i> ^[a] | The Gprof plug-in, which integrates the Gprof performance analysis utility with Eclipse. |
| <i>devtoolset-2-eclipse-manpage</i> ^[a] | The Man Page plug-in, which allows you to view manual pages in Eclipse. |
| <i>devtoolset-2-eclipse-oprofile</i> ^[a] | The OProfile plug-in, which integrates OProfile with Eclipse. |
| <i>devtoolset-2-eclipse-perf</i> ^[a] | The Perf plug-in, which integrates the perf tool with Eclipse. |
| <i>devtoolset-2-eclipse-rpm-editor</i> ^[a] | The Eclipse Spec File Editor, which allows you to maintain RPM spec files. |
| <i>devtoolset-2-eclipse-rpmstubby</i> ^[a] | The RPM Stubby plug-in, which allows you to generate RPM spec files. |
| <i>devtoolset-2-eclipse-systemtap</i> ^[a] | The SystemTap plug-in, which integrates SystemTap with Eclipse. |
| <i>devtoolset-2-eclipse-valgrind</i> ^[a] | The Valgrind plug-in, which integrates Valgrind with Eclipse. |

[a] This package is installed as a dependency of *devtoolset-2-eclipse-linuxtools*.

2.1. Installing Eclipse

In Red Hat Developer Toolset, the Eclipse development environment is provided as a collection of RPM packages and is automatically installed with the *devtoolset-2-ide* package as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#). For a list of available components, see [Table 2.1, “Eclipse Components Included in Red Hat Developer Toolset”](#).

**Note**

The Red Hat Developer Toolset version of Eclipse is only available for Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures. This version fully supports C, C++, and Java development, but does *not* provide support for the Fortran programming language.

2.1.1. Building Eclipse from the Source RPM Packages

It is recommended that you install the Red Hat Developer Toolset version of Eclipse from the official RPM packages distributed by Red Hat. If, for some reason, you need to build Red Hat Developer Toolset 2.1 Eclipse on Red Hat Enterprise Linux 6 from the supplied source RPM (SRPM) packages, complete the following steps:

1. Install the *java-1.7.0-openjdk* package from Red Hat Enterprise Linux 6.3 or newer and all its dependencies.
2. Download, build, and install the *maven* package and all its dependencies from the SRPM packages for Fedora 19.
3. Download, build, and install the *tycho* package and all its dependencies from the SRPM packages for Fedora 19.
4. Download and build all dependent packages listed in the *devtoolset-2-eclipse* SRPM package.
5. Download and build the *devtoolset-2-eclipse* SRPM package.
6. Download and build additional *devtoolset-2-eclipse-** SRPM packages.

2.2. Using Eclipse

To start the Red Hat Developer Toolset version of Eclipse, either select **Applications** → **Programming** → **DTS Eclipse** from the panel, or type the following at a shell prompt:

```
scl enable devtoolset-2 'eclipse'
```

During its startup, Eclipse prompts you to select a *workspace*, that is, a directory in which you want to store your projects. You can either use `~/workspace/`, which is the default option, or click the **Browse** button to browse your file system and select a custom directory. Additionally, you can select the **Use this as the default and do not ask again** check box to prevent Eclipse from displaying this dialog box the next time you run this development environment. When you are done, click the **OK** button to confirm the selection and proceed with the startup.

2.2.1. Using the Red Hat Developer Toolset Toolchain

To use the Red Hat Developer Toolset version of Eclipse with support for the GNU Compiler Collection and binutils from Red Hat Developer Toolset, make sure that the *devtoolset-2-toolchain* package is installed and run the application as described in [Section 2.2, “Using Eclipse”](#). Red Hat Developer Toolset Eclipse uses the Red Hat Developer Toolset toolchain by default.

For detailed instructions on how to install the *devtoolset-2-toolchain* package in your system, see [Section 1.5, “Installing Red Hat Developer Toolset”](#).



Important

If you are working on a project that you previously built with the Red Hat Enterprise Linux version of the GNU Compiler Collection, make sure that you discard all previous build results. To do so, open the project in Eclipse and select **Project** → **Clean** from the menu.

2.2.2. Using the Red Hat Enterprise Linux Toolchain

To use the Red Hat Developer Toolset version of Eclipse with support for the toolchain distributed with Red Hat Enterprise Linux, either uninstall *devtoolset-2-gcc*, *devtoolset-2-binutils*, and related packages from your system, or change the configuration of the project to use absolute paths to the Red Hat Enterprise Linux system versions of **gcc**, **g++**, and **as**.

To uninstall the Red Hat Developer Toolset toolchain from the system, type the following at a shell prompt as **root**:

```
yum remove devtoolset-2-gcc\* devtoolset-2-binutils\*
```

To configure Eclipse to explicitly use the Red Hat Enterprise Linux system versions of the tools for the current project, complete the following steps:

1. In the C/C++ perspective, choose **Project** → **Properties** from the main menu bar to open the project properties.
2. In the menu on the left-hand side of the dialog box, click **C/C++ Build** → **Settings**.
3. Select the **Tool Settings** tab.
4. If you are working on a C project:
 - a. select **GCC C Compiler** or **Cross GCC Compiler** and change the value of the **Command** field to:

```
/usr/bin/gcc
```

- b. select **GCC C Linker** or **Cross GCC Linker** and change the value of the **Command** field to:

```
/usr/bin/gcc
```

- c. select **GCC Assembler** or **Cross GCC Assembler** and change the value of the **Command** field to:

```
/usr/bin/as
```

If you are working on a C++ project:

- a. select **GCC C++ Compiler** or **Cross G++ Compiler** and change the value of the **Command** field to:

```
/usr/bin/g++
```

- b. select **GCC C Compiler** or **Cross GCC Compiler** and change the value of the **Command** field to:

```
/usr/bin/gcc
```

- c. select **GCC C++ Linker** or **Cross G++ Linker** and change the value of the **Command** field to:

```
/usr/bin/g++
```

- d. select **GCC Assembler** or **Cross GCC Assembler** and change the value of the **Command** field to:

```
/usr/bin/as
```

5. Click the **OK** button to save the configuration changes.

2.3. Additional Resources

A detailed description of Eclipse and all its features is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

- Eclipse includes a built-in **Help** system which provides extensive documentation for each integrated feature and tool. This greatly decreases the initial time investment required for new developers to become fluent in its use. The use of this Help section is detailed in the *Red Hat Enterprise Linux Developer Guide* linked below.

Online Documentation

- [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information on Eclipse, including a description of the user interface, overview of available development toolkits, or instructions on how to use it to build RPM packages.

See Also

- [Section A.1, “Changes in Eclipse”](#) provides a comprehensive list of features and improvements over the Eclipse development environment included in Red Hat Enterprise Linux and the previous release of Red Hat Developer Toolset.
- [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 3, GNU Compiler Collection \(GCC\)](#) provides information on how to compile programs written in C, C++, and Fortran on the command line.

Part III. Development Tools

Chapter 3. GNU Compiler Collection (GCC)

The **GNU Compiler Collection**, commonly abbreviated **GCC**, is a portable compiler suite with support for a wide selection of programming languages.

Red Hat Developer Toolset is distributed with **GCC 4.8**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements, including optimization for various new Intel and AMD processors, support for OpenMP 3.1 and link-time optimization. This version also includes experimental support for the C++11 standard, C++11 atomic types, and Transactional Memory. For a detailed list of changes, see [Section B.2, “Changes in GCC”](#).



Note

The Red Hat Developer Toolset version of the GNU Debugger is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures. Notable differences are as follows:

- » On Red Hat Enterprise Linux 6, GCC uses the **STT_GNU_UNIQUE** symbol. On Red Hat Enterprise Linux 5, it does not.
- » On Red Hat Enterprise Linux 6, GCC uses the **-march=i686** command line option by default. On Red Hat Enterprise Linux 5, it uses the **-march=i586** option.
- » On Red Hat Enterprise Linux 6, GCC uses the **-gdwarf-3** command line option by default. On Red Hat Enterprise Linux 5, it uses the **-gdwarf-2** option.
- » The **-fgnu89-inline** command line option has different defaults on Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 5.

3.1. GNU C Compiler

3.1.1. Installing the C Compiler

In Red Hat Developer Toolset, the GNU C compiler is provided by the *devtoolset-2-gcc* package and is automatically installed with *devtoolset-2-toolchain* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

3.1.2. Using the C Compiler

To compile a C program on the command line, run the **gcc** compiler as follows:

```
scl enable devtoolset-2 'gcc -o output_file source_file...'
```

This creates a binary file named **output_file** in the current working directory. If the **-o** option is omitted, the compiler creates a file named **a.out** by default.

When you are working on a project that consists of several source files, it is common to compile an object file for each of the source files first and then link these object files together. This way, when you change a single source file, you can recompile only this file without having to compile the entire project. To compile an object file on the command line, run the following command:

```
scl enable devtoolset-2 'gcc -o object_file -c source_file'
```

This creates an object file named **object_file**. If the **-o** option is omitted, the compiler creates a file

named after the source file with the `.o` file extension. To link object files together and create a binary file, run:

```
scl enable devtoolset-2 'gcc -o output_file object_file...'
```

Note that you can execute any command using the `scl` utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset `gcc` as default:

```
scl enable devtoolset-2 'bash'
```



Note

To verify the version of `gcc` you are using at any point, type the following at a shell prompt:

```
which gcc
```

Red Hat Developer Toolset's `gcc` executable path will begin with `/opt`. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset `gcc`:

```
gcc -v
```



Important

Some newer library features are statically linked into applications built with Red Hat Developer Toolset to support execution on multiple versions of Red Hat Enterprise Linux. This adds a small additional security risk as normal Red Hat Enterprise Linux errata would not change this code. If the need for developers to rebuild their applications due to such an issue arises, Red Hat will signal this via a security erratum. Developers are strongly advised not to statically link their entire application for the same reasons.

Example 3.1. Compiling a C Program on the Command Line

Consider a source file named **hello.c** with the following contents:

```
#include <stdio.h>

int main(int argc, char *argv[]) {
    printf("Hello, World!\n");
    return 0;
}
```

To compile this source code on the command line by using the **gcc** compiler from Red Hat Developer Toolset, type:

```
~]$ scl enable devtoolset-2 'gcc -o hello hello.c'
```

This creates a new binary file called **hello** in the current working directory.

3.1.3. Running a C Program

When **gcc** compiles a program, it creates an executable binary file. To run this program on the command line, change to the directory with the executable file and type:

```
./file_name
```

Example 3.2. Running a C Program on the Command Line

Assuming that you have successfully compiled the **hello** binary file as shown in [Example 3.1, “Compiling a C Program on the Command Line”](#), you can run it by typing the following at a shell prompt:

```
~]$ ./hello
Hello, World!
```

3.2. GNU C++ Compiler**3.2.1. Installing the C++ Compiler**

In Red Hat Developer Toolset, the GNU C++ compiler is provided by the *devtoolset-2-gcc-c++* package and is automatically installed with the *devtoolset-2-toolchain* package as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

3.2.2. Using the C++ Compiler

To compile a C++ program on the command line, run the **g++** compiler as follows:

```
scl enable devtoolset-2 'g++ -o output_file source_file...'
```

This creates a binary file named **output_file** in the current working directory. If the **-o** option is

omitted, the **g++** compiler creates a file named **a.out** by default.

When you are working on a project that consists of several source files, it is common to compile an object file for each of the source files first and then link these object files together. This way, when you change a single source file, you can recompile only this file without having to compile the entire project. To compile an object file on the command line, run the following command:

```
sc1 enable devtoolset-2 'g++ -o object_file -c source_file'
```

This creates an object file named ***object_file***. If the **-o** option is omitted, the **g++** compiler creates a file named after the source file with the **.o** file extension. To link object files together and create a binary file, run:

```
sc1 enable devtoolset-2 'g++ -o output_file object_file...'
```

Note that you can execute any command using the **sc1** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **g++** as default:

```
sc1 enable devtoolset-2 'bash'
```



Note

To verify the version of **g++** you are using at any point, type the following at a shell prompt:

```
which g++
```

Red Hat Developer Toolset's **g++** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **g++**:

```
g++ -v
```



Important

Some newer library features are statically linked into applications built with Red Hat Developer Toolset to support execution on multiple versions of Red Hat Enterprise Linux. This adds a small additional security risk as normal Red Hat Enterprise Linux errata would not change this code. If the need for developers to rebuild their applications due to such an issue arises, Red Hat will signal this via a security erratum. Developers are strongly advised not to statically link their entire application for the same reasons.

Example 3.3. Compiling a C++ Program on the Command Line

Consider a source file named **hello.cpp** with the following contents:

```
#include <iostream>

using namespace std;

int main(int argc, char *argv[]) {
    cout << "Hello, World!" << endl;
    return 0;
}
```

To compile this source code on the command line by using the **g++** compiler from Red Hat Developer Toolset, type:

```
~]$ scl enable devtoolset-2 'g++ -o hello hello.cpp'
```

This creates a new binary file called **hello** in the current working directory.

3.2.3. Running a C++ Program

When **g++** compiles a program, it creates an executable binary file. To run this program on the command line, change to the directory with the executable file and type:

```
./file_name
```

Example 3.4. Running a C++ Program on the Command Line

Assuming that you have successfully compiled the **hello** binary file as shown in [Example 3.3, “Compiling a C++ Program on the Command Line”](#), you can run it by typing the following at a shell prompt:

```
~]$ ./hello
Hello, World!
```

3.3. GNU Fortran Compiler**3.3.1. Installing the Fortran Compiler**

In Red Hat Developer Toolset, the GNU Fortran compiler is provided by the *devtoolset-2-gcc-gfortran* package and is automatically installed with *devtoolset-2-toolchain* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

3.3.2. Using the Fortran Compiler

To compile a Fortran program on the command line, run the **gfortran** compiler as follows:

```
scl enable devtoolset-2 'gfortran -o output_file source_file...'
```

This creates a binary file named **output_file** in the current working directory. If the **-o** option is omitted, the compiler creates a file named **a.out** by default.

When you are working on a project that consists of several source files, it is common to compile an object file for each of the source files first and then link these object files together. This way, when you change a single source file, you can recompile only this file without having to compile the entire project. To compile an object file on the command line, run the following command:

```
scl enable devtoolset-2 'gfortran -o object_file -c source_file'
```

This creates an object file named **object_file**. If the **-o** option is omitted, the compiler creates a file named after the source file with the **.o** file extension. To link object files together and create a binary file, run:

```
scl enable devtoolset-2 'gfortran -o output_file object_file...'
```

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **gfortran** as default:

```
scl enable devtoolset-2 'bash'
```



Note

To verify the version of **gfortran** you are using at any point, type the following at a shell prompt:

```
which gfortran
```

Red Hat Developer Toolset's **gfortran** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **gfortran**:

```
gfortran -v
```



Important

Some newer library features are statically linked into applications built with Red Hat Developer Toolset to support execution on multiple versions of Red Hat Enterprise Linux. This adds a small additional security risk as normal Red Hat Enterprise Linux errata would not change this code. If the need for developers to rebuild their applications due to such an issue arises, Red Hat will signal this via a security erratum. Developers are strongly advised not to statically link their entire application for the same reasons.

Example 3.5. Compiling a Fortran Program on the Command Line

Consider a source file named **hello.f** with the following contents:

```
program hello
  print *, "Hello, World!"
end program hello
```

To compile this source code on the command line by using the **gfortran** compiler from Red Hat Developer Toolset, type:

```
~]$ scl enable devtoolset-2 'gfortran -o hello hello.f'
```

This creates a new binary file called **hello** in the current working directory.

3.3.3. Running a Fortran Program

When **gfortran** compiles a program, it creates an executable binary file. To run this program on the command line, change to the directory with the executable file and type:

```
./file_name
```

Example 3.6. Running a Fortran Program on the Command Line

Assuming that you have successfully compiled the **hello** binary file as shown in [Example 3.5, “Compiling a Fortran Program on the Command Line”](#), you can run it by typing the following at a shell prompt:

```
~]$ ./hello
Hello, World!
```

3.4. Additional Resources

A detailed description of the GNU Compiler Collections and its features is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

- **gcc(1)** — The manual page for the **gcc** compiler provides detailed information on its usage; with few exceptions, **g++** accepts the same command line options as **gcc**. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man gcc'
```

- **gfortran(1)** — The manual page for the **gfortran** compiler provides detailed information on its usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man gfortran'
```

- *C++ Standard Library Documentation* — Documentation on the C++ standard library can be optionally installed by typing the following at a shell prompt as **root**:

```
yum install devtoolset-2-libstdc++-docs
```

Once installed, HTML documentation is available at `/opt/rh/devtoolset-2/root/usr/share/doc/devtoolset-2-libstdc++-docs-4.8.2/html/index.html`.

Online Documentation

- [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides in-depth information about GCC.
- [Using the GNU Compiler Collection](#) — The official GCC manual provides an in-depth description of the GNU compilers and their usage.
- [The GNU C++ Library](#) — The GNU C++ library documentation provides detailed information about the GNU implementation of the standard C++ library.
- [The GNU Fortran Compiler](#) — The GNU Fortran compiler documentation provides detailed information on **gfortran**'s usage.

See Also

- [Section B.2, “Changes in GCC”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of the GNU Compiler Collection and the version distributed in the previous release of Red Hat Developer Toolset, as well as information about the language, ABI, and debugging compatibility.
- [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 2, Eclipse](#) provides a general introduction to the Eclipse development environment, and describes how to use it with the tools from Red Hat Developer Toolset.
- [Chapter 4, binutils](#) explains how to use the binutils, a collection of binary tools to inspect and manipulate object files and binaries.
- [Chapter 5, elfutils](#) explains how to use elfutils, a collection of binary tools to inspect and manipulate ELF files.
- [Chapter 6, dwz](#) explains how to use dwz to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size.
- [Chapter 8, GNU Debugger \(GDB\)](#) provides information on how to debug programs written in C, C++, and Fortran.

Chapter 4. binutils

binutils is a collection of various binary tools such as the GNU linker, GNU assembler, and other utilities that allow you to inspect and manipulate object files and binaries. See [Table 4.1, “Tools Included in binutils for Red Hat Developer Toolset”](#) for a complete list of binary tools that are distributed with the Red Hat Developer Toolset version of binutils.

Red Hat Developer Toolset is distributed with **binutils 2.23.52**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements, including the new **gold** linker, several new command line options, improvements to the linker script language, and support for link-time optimization, compressed debug sections, and new instruction sets. For a detailed list of changes, see [Section B.3, “Changes in binutils”](#).

Table 4.1. Tools Included in binutils for Red Hat Developer Toolset

| Name | Description |
|------------------|---|
| addr2line | Translates addresses into file names and line numbers. |
| ar | Creates, modifies, and extracts files from archives. |
| as | The GNU assembler. |
| c++filt | Decodes mangled C++ symbols. |
| dwp | Combines DWARF object files into a single DWARF package file. |
| elfedit | Examines and edits ELF files. |
| gprof | Display profiling information. |
| ld | The GNU linker. |
| ld.bfd | An alternative to the GNU linker. |
| ld.gold | A new ELF linker. |
| nm | Lists symbols from object files. |
| objcopy | Copies and translates object files. |
| objdump | Displays information from object files. |
| ranlib | Generates an index to the contents of an archive to make access to this archive faster. |
| readelf | Displays information about ELF files. |
| size | Lists section sizes of object or archive files. |
| strings | Displays printable character sequences in files. |
| strip | Discards all symbols from object files. |

4.1. Installing binutils

In Red Hat Developer Toolset, binutils are provided by the *devtoolset-2-binutils* package and are automatically installed with *devtoolset-2-toolchain* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

4.2. Using the GNU Assembler

To produce an object file from an assembly language program, run the **as** tool as follows:

```
scl enable devtoolset-2 'as [option...] -o object_file source_file'
```

This creates an object file named ***object_file*** in the current working directory.

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **as** as default:

```
scl enable devtoolset-2 'bash'
```



Note

To verify the version of **as** you are using at any point, type the following at a shell prompt:

```
which as
```

Red Hat Developer Toolset's **as** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **as**:

```
as -v
```

4.3. Using the GNU Linker

To create an executable binary file or a library from object files, run the **ld** tool as follows:

```
scl enable devtoolset-2 'ld [option...] -o output_file object_file...'
```

This creates a binary file named ***output_file*** in the current working directory. If the **-o** option is omitted, the compiler creates a file named **a.out** by default.

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **ld** as default:

```
scl enable devtoolset-2 'bash'
```

**Note**

To verify the version of **ld** you are using at any point, type the following at a shell prompt:

```
which ld
```

Red Hat Developer Toolset's **ld** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **ld**:

```
ld -v
```

4.4. Using Other Binary Tools

The binutils provide many binary tools other than a linker and assembler. For a complete list of these tools, see [Table 4.1, “Tools Included in binutils for Red Hat Developer Toolset”](#).

To execute any of the tools that are part of binutils, run the command as follows:

```
scl enable devtoolset-2 'tool [option...] file_name'
```

See [Table 4.1, “Tools Included in binutils for Red Hat Developer Toolset”](#) for a list of tools that are distributed with binutils. For example, to use the **objdump** tool to inspect an object file, type:

```
scl enable devtoolset-2 'objdump [option...] object_file'
```

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset binary tools as default:

```
scl enable devtoolset-2 'bash'
```

**Note**

To verify the version of binutils you are using at any point, type the following at a shell prompt:

```
which objdump
```

Red Hat Developer Toolset's **objdump** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **objdump**:

```
objdump -v
```

4.5. Additional Resources

A detailed description of binutils is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

- **as(1), ld(1), addr2line(1), ar(1), c++filt(1), dwp(1), elfedit(1), gprof(1), nm(1), objcopy(1), objdump(1), ranlib(1), readelf(1), size(1), strings(1), strip(1)**, — Manual pages for various binutils tools provide more information about their respective usage. To display a manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man tool'
```

Online Documentation

- [Documentation for binutils](#) — The binutils documentation provides an in-depth description of the binary tools and their usage.

See Also

- [Section B.3, “Changes in binutils”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of binutils and the version distributed in the previous release of Red Hat Developer Toolset.
- [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 5, elfutils](#) explains how to use elfutils, a collection of binary tools to inspect and manipulate ELF files.
- [Chapter 3, GNU Compiler Collection \(GCC\)](#) provides information on how to compile programs written in C, C++, and Fortran.

Chapter 5. elfutils

elfutils is a collection of various binary tools such as **eu-objdump**, **eu-readelf**, and other utilities that allow you to inspect and manipulate ELF files. See [Table 5.1, “Tools Included in elfutils for Red Hat Developer Toolset”](#) for a complete list of binary tools that are distributed with the Red Hat Developer Toolset version of elfutils.

Red Hat Developer Toolset is distributed with **elfutils 0.157**. This version is more recent than the version included in Red Hat Enterprise Linux and the previous release of Red Hat Developer Toolset, and provides numerous bug fixes and enhancements.

Table 5.1. Tools Included in elfutils for Red Hat Developer Toolset

| Name | Description |
|------------------------------|--|
| eu-addr2line | Translates addresses into file names and line numbers. |
| eu-ar | Creates, modifies, and extracts files from archives. |
| eu-elfcmp | Compares relevant parts of two ELF files for equality. |
| eu-elflint | Verifies that ELF files are compliant with the <i>generic ABI</i> (gABI) and <i>processor-specific supplement ABI</i> (psABI) specification. |
| eu-findtextrel | Locates the source of text relocations in files. |
| eu-make-debug-archive | Creates an offline archive for debugging. |
| eu-nm | Lists symbols from object files. |
| eu-objdump | Displays information from object files. |
| eu-ranlib | Generates an index to the contents of an archive to make access to this archive faster. |
| eu-readelf | Displays information about ELF files. |
| eu-size | Lists section sizes of object or archive files. |
| eu-strings | Displays printable character sequences in files. |
| eu-strip | Discards all symbols from object files. |
| eu-unstrip | Combines stripped files with separate symbols and debug information. |

5.1. Installing elfutils

In Red Hat Developer Toolset, elfutils is provided by the *devtoolset-2-elfutils* package and is automatically installed with *devtoolset-2-toolchain* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).



Note

The Red Hat Developer Toolset version of elfutils is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures.

5.2. Using elfutils

To execute any of the tools that are part of elfutils, run the command as follows:

```
scl enable devtoolset-2 'tool [option...] file_name'
```

See [Table 5.1, “Tools Included in elfutils for Red Hat Developer Toolset”](#) for a list of tools that are distributed with elfutils. For example, to use the **eu-objdump** tool to inspect an object file, type:

```
scl enable devtoolset-2 'eu-objdump [option...] object_file'
```

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset binary tools as default:

```
scl enable devtoolset-2 'bash'
```



Note

To verify the version of elfutils you are using at any point, type the following at a shell prompt:

```
which eu-objdump
```

Red Hat Developer Toolset's **eu-objdump** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **eu-objdump**:

```
eu-objdump -V
```

5.3. Additional Resources

A detailed description of elfutils is beyond the scope of this book. For more information, see the resources listed below.

See Also

- [Section A.3, “Changes in elfutils”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of elfutils and the version distributed in the previous release of Red Hat Developer Toolset.
- [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 3, GNU Compiler Collection \(GCC\)](#) provides information on how to compile programs written in C, C++, and Fortran.
- [Chapter 4, binutils](#) explains how to use the binutils, a collection of binary tools to inspect and manipulate object files and binaries.
- [Chapter 6, dwz](#) explains how to use dwz to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size.

Chapter 6. dwz

dwz is a command line tool that attempts to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size. To do so, **dwz** replaces DWARF information representation with equivalent smaller representation where possible, and reduces the amount of duplication by using techniques from *Appendix E* of the *DWARF Standard*.

Red Hat Developer Toolset is distributed with **dwz 0.11**.

6.1. Installing dwz

In Red Hat Developer Toolset, the **dwz** utility is provided by the `devtoolset-2-dwz` package and is automatically installed with `devtoolset-2-toolchain` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).



Note

The Red Hat Developer Toolset version of **dwz** is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures.

6.2. Using dwz

To optimize DWARF debugging information in a binary file, run the **dwz** tool as follows:

```
sc1 enable devtoolset-2 'dwz [option...] file_name'
```

Note that you can execute any command using the **sc1** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **dwz** as default:

```
sc1 enable devtoolset-2 'bash'
```



Note

To verify the version of **dwz** you are using at any point, type the following at a shell prompt:

```
which dwz
```

Red Hat Developer Toolset's **dwz** executable path will begin with `/opt`. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **dwz**:

```
dwz -v
```

6.3. Additional Resources

A detailed description of **dwz** and its features is beyond the scope of this book. For more information,

see the resources listed below.

Installed Documentation

- **dwz(1)** — The manual page for the **dwz** utility provides detailed information on its usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man dwz'
```

See Also

- [Section B.5, “Changes in dwz”](#) provides a comprehensive list of features and enhancements over the version of **dwz** distributed in the previous release of Red Hat Developer Toolset.
- [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 3, GNU Compiler Collection \(GCC\)](#) provides information on how to compile programs written in C, C++, and Fortran.
- [Chapter 4, binutils](#) explains how to use the binutils, a collection of binary tools to inspect and manipulate object files and binaries.
- [Chapter 5, elfutils](#) explains how to use elfutils, a collection of binary tools to inspect and manipulate ELF files.

Part IV. Revision Control Systems

Chapter 7. Git

Git is a *distributed revision control system* with a peer-to-peer architecture. As opposed to centralized version control systems with a client-server model, Git ensures that each working copy of a Git repository is its exact copy with complete revision history. This not only allows you to work on and contribute to projects without the need to have permission to push your changes to their official repositories, but also makes it possible for you to work with no network connection.

Red Hat Developer Toolset is distributed with **Git 1.8.4**. This version is more recent than the version included in Red Hat Enterprise Linux and provides a number of bug fixes and enhancements. For a detailed list of changes, see [Section A.4, “Changes in Git”](#).

7.1. Installing Git

In Red Hat Developer Toolset, the **Git** revision control system is provided as a collection of RPM packages and is automatically installed with the *devtoolset-2-vc* package as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).



Note

The Red Hat Developer Toolset version of Git is only available for Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures.

7.2. Configuring Git

7.2.1. Configuring the Default Text Editor

Certain Git commands such as **git commit** require the user to write a short message or make some changes in an external text editor. To determine which text editor to start, Git attempts to read the value of the **GIT_EDITOR** environment variable, the **core.editor** configuration option, the **VISUAL** environment variable, and finally the **EDITOR** environment variable in this particular order. If none of these options and variables are specified, the **git** command starts **vi** as a reasonable default option.

To change the value of the **core.editor** configuration option in order to specify a different text editor, type the following at a shell prompt:

```
scl enable devtoolset-2 'git config --global core.editor command'
```

Replace **command** with the command to be used to start the selected text editor.

Example 7.1. Configuring the Default Text Editor

To configure Git to use **vim** as the default text editor, type the following at a shell prompt:

```
~]$ scl enable devtoolset-2 'git config --global core.editor vim'
```

7.2.2. Setting Up User Information

In Git, each commit (or revision) is associated with the full name and email of the person responsible for

it. By default, Git uses an identity based on the user name and the host name.

To change the full name associated with your Git commits, type the following at a shell prompt:

```
scl enable devtoolset-2 'git config --global user.name "full name"'
```

To change the email address associated with your Git commits, type:

```
scl enable devtoolset-2 'git config --global user.email "email_address"'
```

Example 7.2. Setting Up User Information

To configure Git to use **John Doe** as your full name and **john@example.com** as your email address, type the following at a shell prompt:

```
~]$ scl enable devtoolset-2 'git config --global user.name "John Doe" '
~]$ scl enable devtoolset-2 'git config --global user.email
"john@example.com" '
```

7.3. Using Git

7.3.1. Creating a New Repository

A *repository* is a place where Git stores all files that are under revision control, as well as additional data related to these files, such as the complete history of changes or information about who made those changes and when. Unlike in centralized revision control systems like Subversion or CVS, a Git repository and a *working directory* are usually the same. A typical Git repository also only stores a single project and when publicly accessible, it allows anyone to create its clone with a complete revision history.

Initializing an Empty Repository

To create a new, empty Git repository, change to the directory in which you want to keep the repository and type the following at a shell prompt:

```
scl enable devtoolset-2 'git init'
```

This creates a hidden directory named **.git** in which all repository information is stored.

Importing Data to a Repository

To put an existing project under revision control, create a Git repository in the directory with the project and run the following command:

```
scl enable devtoolset-2 'git add .'
```

This marks all files and directories in the current working directory as ready to be added to the Git repository. To proceed and actually add this content to the repository, commit the changes by typing the following at a shell prompt:

```
scl enable devtoolset-2 'git commit [-m "commit message"]'
```

Replace **commit message** with a short description of your revision. If you omit the **-m** option, this

command allows you to write the commit message in an external text editor. For information on how to configure the default text editor, see [Section 7.2.1, “Configuring the Default Text Editor”](#).

7.3.2. Cloning an Existing Repository

Cloning a Git Repository

To clone an existing Git repository, type the following at a shell prompt:

```
scl enable devtoolset-2 'git clone git_repository [directory]'
```

Replace ***git_repository*** with a URL or a path to the Git repository you want to clone, and ***directory*** with a path to the directory in which you want to store the clone.

Cloning a Subversion Repository

When the *devtoolset-2-git-svn* package is installed in the system, Git allows you to create a Git clone of a Subversion repository. To do so, type the following at a shell prompt:

```
scl enable devtoolset-2 'git svn clone svn_repository [directory]'
```

Replace ***svn_repository*** with a URL or a path to the Subversion repository you want to clone, and ***directory*** with a path to the directory in which you want to store the clone.

For information on how to install the *devtoolset-2-git-svn* package, see [Section 1.5.3, “Installing Optional Packages”](#).

Cloning a CVS Repository

When the *devtoolset-2-git-cvs* package is installed in the system, Git allows you to create a Git clone of a CVS repository. To do so, change to the directory with your working copy of the CVS repository and type the following at a shell prompt:

```
scl enable devtoolset-2 'git cvsimport -C directory cvs_module'
```

Replace ***cvs_module*** with the name of the subdirectory in which the project is stored, and ***directory*** with a path to the directory in which you want to store the clone.

For information on how to install the *devtoolset-2-git-cvs* package, see [Section 1.5.3, “Installing Optional Packages”](#).

7.3.3. Adding, Renaming, and Deleting Files

Adding Files and Directories

To add an existing file to a Git repository and put it under revision control, change to the directory with your local Git repository and type the following at a shell prompt:

```
scl enable devtoolset-2 'git add file...'
```

Replace ***file*** with the file or files you want to add. This command marks the selected file or files as ready to be added to the Git repository. Similarly, to add all files that are stored in a certain directory to a Git repository, type:

```
scl enable devtoolset-2 'git add directory...'
```

Replace **directory** with the directory or directories you want to add. This command marks all files in the selected directory or directories as ready to be added to the Git repository.

To proceed and actually add this content to the repository, commit the changes as described in [Section 7.3.5, “Committing Changes”](#).

Renaming Files and Directories

To rename an existing file or directory in a Git repository, change to the directory with your local Git repository and type the following at a shell prompt:

```
scl enable devtoolset-2 'git mv old_name new_name'
```

Replace **old_name** with the current name of the file or directory and **new_name** with the new name. This command renames the selected file or directory and marks it as ready to be renamed in the Git repository.

To proceed and actually rename the content in the repository, commit the changes as described in [Section 7.3.5, “Committing Changes”](#).

Deleting Files and Directories

To delete an existing file from a Git repository, change to the directory with your local Git repository and type the following at a shell prompt:

```
scl enable devtoolset-2 'git rm file...'
```

Replace **file** with the file or files you want to delete. This command deletes all selected files and marks them as ready to be deleted from the Git repository. Similarly, to delete all files that are stored in a certain directory from a Git repository, type:

```
scl enable devtoolset-2 'git rm -r directory...'
```

Replace **directory** with the directory or directories you want to delete. This command deletes all selected directories and marks them as ready to be deleted from the Git repository.

To proceed and actually delete this content from the repository, commit the changes as described in [Section 7.3.5, “Committing Changes”](#).

7.3.4. Viewing Changes

Viewing the Current Status

To determine the current status of your local Git repository, change to the directory with the repository and type the following command at a shell prompt:

```
scl enable devtoolset-2 'git status'
```

This command displays information about all uncommitted changes in the repository (**new file**, **renamed**, **deleted**, or **modified**) and tells you which changes will be applied the next time you commit them. For information on how to commit your changes, see [Section 7.3.5, “Committing Changes”](#).

Viewing Differences

To view all changes in a Git repository, change to the directory with the repository and type the following

at a shell prompt:

```
scl enable devtoolset-2 'git diff'
```

This command displays changes between the files in the repository and their latest revision. If you are only interested in changes in a particular file, supply its name on the command line as follows:

```
scl enable devtoolset-2 'git diff file...'
```

Replace *file* with the file or files you want to view.

7.3.5. Committing Changes

To apply your changes to a Git repository and create a new revision, change to the directory with the repository and type the following command at a shell prompt:

```
scl enable devtoolset-2 'git commit [-m "commit message"]'
```

Replace *commit message* with a short description of your revision. This command commits all changes in files that are explicitly marked as ready to be committed. To commit changes in all files that are under revision control, add the **-a** command line option as follows:

```
scl enable devtoolset-2 'git commit -a [-m "commit message"]'
```

Note that if you omit the **-m** option, the command allows you to write the commit message in an external text editor. For information on how to configure the default text editor, see [Section 7.2.1, “Configuring the Default Text Editor”](#).

7.3.6. Sharing Changes

Unlike in centralized version control systems such as CVS or Subversion, when working with Git, project contributors usually do not make their changes in a single, central repository. Instead, they either create a publicly accessible clone of their local repository, or submit their changes to others over email as patches.

Pushing Changes to a Public Git Repository

To push your changes to a publicly accessible Git repository, change to the directory with your local repository and type the following at a shell prompt:

```
scl enable devtoolset-2 'git push remote_repository'
```

Replace *remote_repository* with the name of the remote repository you want to push your changes to. Note that the repository from which you originally cloned your local copy is automatically named **origin**.

Creating Patches from Individual Commits

To create patches from your commits, change to the directory with your local Git repository and type the following at a shell prompt:

```
scl enable devtoolset-2 'git format-patch remote_repository'
```

Replace *remote_repository* with the name of the remote repository from which you made your local copy. This creates a patch for each commit that is not present in this remote repository.

7.3.7. Updating a Git Repository

To update your local copy of a Git repository and get the latest changes from a remote repository, change to the directory with your local Git repository and type the following at a shell prompt:

```
scl enable devtoolset-2 'git fetch remote_repository'
```

Replace *remote_repository* with the name of the remote repository. This command fetches information about the current status of the remote repository, allowing you to review these changes before applying them to your local copy. To proceed and merge these changes with what you have in your local Git repository, type:

```
scl enable devtoolset-2 'git merge remote_repository'
```

Alternatively, you can perform both these steps at the same time by using the following command instead:

```
scl enable devtoolset-2 'git pull remote_repository'
```

7.4. Additional Resources

A detailed description of Git and its features is beyond the scope of this book. For more information about this revision control system, refer to the resources listed below.

Installed Documentation

- **gittutorial(7)** — The manual page named **gittutorial** provides a brief introduction to Git and its usage. To display the manual page for the version of Git included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man gittutorial'
```

- **gittutorial-2(7)** — The manual page named **gittutorial-2** provides the second part of a brief introduction to Git and its usage. To display the manual page for the version of Git included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man gittutorial-2'
```

- *Git User's Manual* — HTML documentation for Git is located at `/opt/rh/devtoolset-2/root/usr/share/doc/devtoolset-2-git-1.8.4/user-manual.html`.

Online Documentation

- [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information about Git and revision control systems in general.
- [Pro Git](#) — The online version of the *Pro Git* book provides a detailed description of Git, its concepts and its usage.

See Also

- [Section A.4, “Changes in Git”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of Git.
- [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more

information on how to install it on your system.

Part V. Debugging Tools

Chapter 8. GNU Debugger (GDB)

The **GNU Debugger**, commonly abbreviated as **GDB**, is a command line tool that can be used to debug programs written in various programming languages. It allows you to inspect memory within the code being debugged, control the execution state of the code, detect the execution of particular sections of code, and much more.

Red Hat Developer Toolset is distributed with **GDB 7.6.1**. This version is more recent than the version included in Red Hat Enterprise Linux and the previous release of Red Hat Developer Toolset, and provides numerous bug fixes and enhancements, including improved support for Python scripting, ambiguous line specifications, and tracepoints, as well as improved inferior control commands and handling of C++ debuggee executables. For a detailed list of changes, see [Section A.5, “Changes in GDB”](#).

8.1. Installing the GNU Debugger

In Red Hat Developer Toolset, the GNU Debugger is provided by the *devtoolset-2-gdb* package and is automatically installed with *devtoolset-2-toolchain* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).



Note

The Red Hat Developer Toolset version of the GNU Debugger is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures. On Red Hat Enterprise Linux 5, the GNU Debugger does not support build IDs and as a consequence, it is unable to determine the versions of binaries that created a particular core dump file. In addition, the GNU Debugger for Red Hat Enterprise Linux 5 does not advise users which RPM packages with debugging information to install.

8.2. Preparing a Program for Debugging

Compiling Programs with Debugging Information

To compile a C program with debugging information that can be read by the GNU Debugger, make sure the **gcc** compiler is run with the **-g** option. To do so on the command line, use a command in the following form:

```
scl enable devtoolset-2 'gcc -g -o output_file input_file...'
```

Similarly, to compile a C++ program with debugging information, run:

```
scl enable devtoolset-2 'g++ -g -o output_file input_file...'
```

Example 8.1. Compiling a C Program With Debugging Information

Consider a source file named **fibonacci.c** that has the following contents:

```
#include <stdio.h>
#include <limits.h>

int main (int argc, char *argv[]) {
    unsigned long int a = 0;
    unsigned long int b = 1;
    unsigned long int sum;

    while (b < LONG_MAX) {
        printf("%ld ", b);
        sum = a + b;
        a = b;
        b = sum;
    }

    return 0;
}
```

To compile this program on the command line using GCC from Red Hat Developer Toolset with debugging information for the GNU Debugger, type:

```
~]$ scl enable devtoolset-2 'gcc -g -o fibonacci fibonacci.c'
```

This creates a new binary file called **fibonacci** in the current working directory.

Installing Debugging Information for Existing Packages

To install debugging information for a package that is already installed on the system, type the following at a shell prompt as **root**:

```
debuginfo-install package_name
```

Note that the *yum-utils* package must be installed for the **debuginfo-install** utility to be available on your system.

Example 8.2. Installing Debugging Information for the glibc Package

To install debugging information for the *glibc* package, type:

```
~]# debuginfo-install glibc
Loaded plugins: product-id, refresh-packagekit, subscription-manager
--> Running transaction check
---> Package glibc-debuginfo.x86_64 0:2.12-1.47.el6_2.5 will be installed
...
```

8.3. Running the GNU Debugger

To run the GNU Debugger on a program you want to debug, type the following at a shell prompt:

```
scl enable devtoolset-2 'gdb file_name'
```

This starts the **gdb** debugger in interactive mode and displays the default prompt, **(gdb)**. To quit the debugging session and return to the shell prompt, run the following command at any time:

```
quit
```

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **gdb** as default:

```
scl enable devtoolset-2 'bash'
```



Note

To verify the version of **gdb** you are using at any point, type the following at a shell prompt:

```
which gdb
```

Red Hat Developer Toolset's **gdb** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **gdb**:

```
gdb -v
```

Example 8.3. Running the gdb Utility on the fibonacci Binary File

Assuming that you have successfully compiled the **fibonacci** binary file as shown in [Example 8.1, "Compiling a C Program With Debugging Information"](#), you can start debugging it with **gdb** by typing the following at a shell prompt:

```
~]$ scl enable devtoolset-2 'gdb fibonacci'
GNU gdb (GDB) Red Hat Enterprise Linux (7.4.50.20120120-43.el6)
Copyright (C) 2012 Free Software Foundation, Inc.
License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law. Type "show copying"
and "show warranty" for details.
This GDB was configured as "x86_64-redhat-linux-gnu".
For bug reporting instructions, please see:
<http://www.gnu.org/software/gdb/bugs/>.
(gdb)
```

8.4. Listing Source Code

To view the source code of the program you are debugging, run the following command:

list

Before you start the execution of the program you are debugging, **gdb** displays first ten lines of the source code and any subsequent use of this command lists another ten lines. Once you start the execution, **gdb** displays the lines that are surrounding the line on which the execution stops, typically when you set a breakpoint.

You can also display the code that is surrounding a particular line. To do so, run the command in the following form:

```
list [file_name:]line_number
```

Similarly, to display the code that is surrounding the beginning of a particular function, run:

```
list [file_name:]function_name
```

Note that you can change the number of lines the **list** command displays by running the following command:

```
set listsize number
```

Example 8.4. Listing the Source Code of the fibonacci Binary File

The **fibonacci.c** file listed in [Example 8.1, “Compiling a C Program With Debugging Information”](#) has exactly 17 lines. Assuming that you have compiled it with debugging information and you want the **gdb** utility to be capable of listing the entire source code, you can run the following command to change the number of listed lines to 20:

```
(gdb) set listsize 20
```

You can now display the entire source code of the file you are debugging by running the **list** command with no additional arguments:

```
(gdb) list
1      #include <stdio.h>
2      #include <limits.h>
3
4      int main (int argc, char *argv[]) {
5          unsigned long int a = 0;
6          unsigned long int b = 1;
7          unsigned long int sum;
8
9          while (b < LONG_MAX) {
10             printf("%ld ", b);
11             sum = a + b;
12             a = b;
13             b = sum;
14         }
15
16         return 0;
17     }
```

8.5. Setting Breakpoints

Setting a New Breakpoint

To set a new breakpoint at a certain line, run the following command:

```
break [file_name:]line_number
```

You can also set a breakpoint on a certain function:

```
break [file_name:]function_name
```

Example 8.5. Setting a New Breakpoint

Assuming that you have compiled the `fibonacci.c` file listed in [Example 8.1, “Compiling a C Program With Debugging Information”](#) with debugging information, you can set a new breakpoint at line 10 by running the following command:

```
(gdb) break 10
Breakpoint 1 at 0x4004e5: file fibonacci.c, line 10.
```

Listing Breakpoints

To display a list of currently set breakpoints, run the following command:

```
info breakpoints
```

Example 8.6. Listing Breakpoints

Assuming that you have followed the instructions in [Example 8.5, “Setting a New Breakpoint”](#), you can display the list of currently set breakpoints by running the following command:

```
(gdb) info breakpoints
Num      Type           Disp Enb Address            What
1        breakpoint     keep y   0x00000000004004e5 in main at fibonacci.c:10
```

Deleting Existing Breakpoints

To delete a breakpoint that is set at a certain line, run the following command:

```
clear line_number
```

Similarly, to delete a breakpoint that is set on a certain function, run:

```
clear function_name
```

You can also delete all breakpoints at once. To do so, run the `clear` command with no additional arguments:

```
clear
```

Example 8.7. Deleting an Existing Breakpoint

Assuming that you have compiled the `fibonacci.c` file listed in [Example 8.1, “Compiling a C Program With Debugging Information”](#) with debugging information, you can set a new breakpoint at line 7 by running the following command:

```
(gdb) break 7
Breakpoint 2 at 0x4004e3: file fibonacci.c, line 7.
```

To remove this breakpoint, type:

```
(gdb) clear 7
Deleted breakpoint 2
```

8.6. Starting Execution

To start execution of the program you are debugging, run the following command:

```
run
```

If the program accepts any command line arguments, you can provide them as arguments to the **run** command:

```
run argument...
```

The execution stops when a first breakpoint (if any) is reached, when an error occurs, or when the program terminates.

Example 8.8. Executing the fibonacci Binary File

Assuming that you have followed the instructions in [Example 8.5, “Setting a New Breakpoint”](#), you can execute the `fibonacci` binary file by running the following command:

```
(gdb) run
Starting program: /home/john/fibonacci

Breakpoint 1, main (argc=1, argv=0x7fffffffe4d8) at fibonacci.c:10
10      printf("%ld ", b);
```

8.7. Displaying Current Values

The **gdb** utility allows you to display the value of almost anything that is relevant to the program, from a variable of any complexity to a valid expression or even a library function. However, the most common task is to display the value of a variable.

To display the current value of a certain variable, run the following command:

```
print variable_name
```

Example 8.9. Displaying the Current Values of Variables

Assuming that you have followed the instructions in [Example 8.8, “Executing the fibonacci Binary File”](#) and the execution of the **fibonacci** binary stopped after reaching the breakpoint at line 10, you can display the current values of variables **a** and **b** as follows:

```
(gdb) print a
$1 = 0
(gdb) print b
$2 = 1
```

8.8. Continuing Execution

To resume the execution of the program you are debugging after it reached a breakpoint, run the following command:

```
continue
```

The execution stops again when another breakpoint is reached. To skip a certain number of breakpoints (typically when you are debugging a loop), you can run the **continue** command in the following form:

```
continue number
```

The **gdb** utility also allows you to stop the execution after executing a single line of code. To do so, run:

```
step
```

Finally, you can execute a certain number of lines by using the **step** command in the following form:

```
step number
```

Example 8.10. Continuing the Execution of the fibonacci Binary File

Assuming that you have followed the instructions in [Example 8.8, “Executing the fibonacci Binary File”](#) and the execution of the `fibonacci` binary stopped after reaching the breakpoint at line 10, you can resume the execution by running the following command:

```
(gdb) continue
Continuing.

Breakpoint 1, main (argc=1, argv=0x7fffffffe4d8) at fibonacci.c:10
10      printf("%ld ", b);
```

The execution stops the next time the breakpoint is reached. To execute next three lines of code, type:

```
(gdb) step 3
13      b = sum;
```

This allows you to verify the current value of the `sum` variable before it is assigned to `b`:

```
(gdb) print sum
$3 = 2
```

8.9. Additional Resources

A detailed description of the GNU Debugger and all its features is beyond the scope of this book. For more information, see the resources listed below.

Online Documentation

- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information on the GNU Debugger and debugging.
- ▶ [GDB Documentation](#) — The official GDB documentation includes the *GDB User Manual* and other reference material.

See Also

- ▶ [Section A.5, “Changes in GDB”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of the GNU Debugger and the version distributed in the previous release of Red Hat Developer Toolset.
- ▶ [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 2, Eclipse](#) provides a general introduction to the Eclipse development environment, and describes how to use it with the tools from Red Hat Developer Toolset.
- ▶ [Chapter 3, GNU Compiler Collection \(GCC\)](#) provides further information on how to compile programs written in C, C++, and Fortran.
- ▶ [Chapter 9, strace](#) documents how to use the `strace` utility to monitor system calls that a program uses and signals it receives.
- ▶ [Chapter 10, memstomp](#) documents how to use the `memstomp` utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.

Chapter 9. strace

strace is a diagnostic and debugging tool for the command line that can be used to trace system calls that are made and received by a running process. It records the name of each system call, its arguments, and its return value, as well as signals received by the process and other interactions with the kernel, and prints this record to standard error output or a selected file.

Red Hat Developer Toolset is distributed with **strace** 4.7.

9.1. Installing strace

In Red Hat Enterprise Linux, the **strace** utility is provided by the *devtoolset-2-strace* package and is automatically installed with *devtoolset-2-toolchain* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).



Note

The Red Hat Developer Toolset version of **strace** is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures.

9.2. Using strace

To run the **strace** utility on a program you want to analyze, type the following at a shell prompt:

```
sc1 enable devtoolset-2 'strace program [argument...]
```

Replace **program** with the name of the program you want to analyze, and **argument** with any command line options and arguments you want to supply to this program. Alternatively, you can run the utility on an already running process by using the **-p** command line option followed by the process ID:

```
sc1 enable devtoolset-2 'strace -p process_id'
```

Note that you can execute any command using the **sc1** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **strace** as default:

```
sc1 enable devtoolset-2 'bash'
```



Note

To verify the version of **strace** you are using at any point, type the following at a shell prompt:

```
which strace
```

Red Hat Developer Toolset's **strace** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **strace**:

```
strace -V
```

9.2.1. Redirecting Output to a File

By default, **strace** prints the name of each system call, its arguments and the return value to standard error output. To redirect this output to a file, use the **-o** command line option followed by the file name:

```
scl enable devtoolset-2 'strace -o file_name program [argument...]'
```

Replace **file_name** with the name of the file.

Example 9.1. Redirecting Output to a File

Consider a slightly modified version of the **fibonacci** file from [Example 8.1, “Compiling a C Program With Debugging Information”](#). This executable file displays the Fibonacci sequence and optionally allows you to specify how many members of this sequence to list. To run the **strace** utility on this file and redirect the trace output to **fibonacci.log**, type:

```
~]$ scl enable devtoolset-2 'strace -o fibonacci.log ./fibonacci 20'
1 1 2 3 5 8 13 21 34 55 89 144 233 377 610 987 1597 2584 4181 6765
```

This creates a new plain-text file called **fibonacci.log** in the current working directory.

9.2.2. Tracing Selected System Calls

To trace only a selected set of system calls, run the **strace** utility with the **-e** command line option:

```
scl enable devtoolset-2 'strace -e expression program [argument...]'
```

Replace **expression** with a comma-separated list of system calls to trace or any of the keywords listed in [Table 9.1, “Commonly Used Values of the -e Option”](#). For a detailed description of all available values, see the **strace(1)** manual page.

Table 9.1. Commonly Used Values of the -e Option

| Value | Description |
|----------------|---|
| file | System calls that accept a file name as an argument. |
| process | System calls that are related to process management. |
| network | System calls that are related to networking. |
| signal | System calls that are related to signal management. |
| ipc | System calls that are related to inter-process communication (IPC). |
| desc | System calls that are related to file descriptors. |

Example 9.2. Tracing Selected System Calls

Consider the **employee** file from [Example 10.1, “Using memstomp”](#). To run the **strace** utility on this executable file and trace only the **mmap** and **munmap** system calls, type:

```
~]$ scl enable devtoolset-2 'strace -e mmap,munmap ./employee'
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0x7f896c744000
mmap(NULL, 61239, PROT_READ, MAP_PRIVATE, 3, 0) = 0x7f896c735000
mmap(0x3146a00000, 3745960, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE, 3,
0) = 0x3146a00000
mmap(0x3146d89000, 20480, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x189000) = 0x3146d89000
mmap(0x3146d8e000, 18600, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0x3146d8e000
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0x7f896c734000
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0x7f896c733000
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0x7f896c732000
munmap(0x7f896c735000, 61239) = 0
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0x7f896c743000
John,john@example.comDoe,
+++ exited with 0 +++
```

9.2.3. Displaying Time Stamps

To prefix each line of the trace with the exact time of the day in hours, minutes, and seconds, run the **strace** utility with the **-t** command line option:

```
scl enable devtoolset-2 'strace -t program [argument...]'
```

To also display milliseconds, supply the **-t** option twice:

```
scl enable devtoolset-2 'strace -tt program [argument...]'
```

To prefix each line of the trace with the time required to execute the respective system call, use the **-r** command line option:

```
scl enable devtoolset-2 'strace -r program [argument...]
```

Example 9.3. Displaying Time Stamps

Consider an executable file named **pwd**. To run the **strace** utility on this file and include time stamps in the output, type:

```
~]$ scl enable devtoolset-2 'strace -tt ./pwd'
19:43:28.011815 execve("./pwd", ["/usr/bin/pwd"], [/* 36 vars */]) = 0
19:43:28.012128 brk(0) = 0xcd3000
19:43:28.012174 mmap(NULL, 4096, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x7fc869cb0000
19:43:28.012427 open("/etc/ld.so.cache", O_RDONLY) = 3
19:43:28.012446 fstat(3, {st_mode=S_IFREG|0644, st_size=61239, ...}) = 0
19:43:28.012464 mmap(NULL, 61239, PROT_READ, MAP_PRIVATE, 3, 0) =
0x7fc869ca1000
19:43:28.012483 close(3) = 0
...
19:43:28.013410 +++ exited with 0 +++
```

9.2.4. Displaying a Summary

To display a summary of how much time was required to execute each system call, how many times were these system calls executed, and how many errors were encountered during their execution, run the **strace** utility with the **-c** command line option:

```
scl enable devtoolset-2 'strace -c program [argument...]
```

Example 9.4. Displaying a Summary

Consider an executable file named **lsblk**. To run the **strace** utility on this file and display a trace summary, type:

```
~]$ scl enable devtoolset-2 'strace -c ./lsblk > /dev/null'
% time      seconds  usecs/call   calls    errors syscall
-----
 80.88      0.000055         1       106        16 open
 19.12      0.000013         0       140         0 munmap
  0.00      0.000000         0       148         0 read
  0.00      0.000000         0         1         0 write
  0.00      0.000000         0      258         0 close
  0.00      0.000000         0        37         2 stat
...
-----
100.00      0.000068                1790        35 total
```

9.3. Additional Resources

A detailed description of **strace** and its features is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

- **strace(1)** — The manual page for the **strace** utility provides detailed information about its usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man strace'
```

See Also

- [Section B.7, “Changes in strace”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of strace.
- [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 8, GNU Debugger \(GDB\)](#) provides information on how to debug programs written in C, C++, and Fortran.
- [Chapter 10, memstomp](#) documents how to use the memstomp utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.

Chapter 10. memstomp

memstomp is a command line tool that can be used to identify function calls with overlapping memory regions in situations when such an overlap is not permitted by various standards. It intercepts calls to the library functions listed in [Table 10.1, “Function Calls Inspected by memstomp”](#) and for each memory overlap, it displays a detailed backtrace to help you debug the problem.

Similarly to **Valgrind**, the **memstomp** utility inspects applications without the need to recompile them. However, it is much faster than this tool and therefore serves as a convenient alternative to it.

Red Hat Developer Toolset is distributed with **memstomp 0.1.4**.

Table 10.1. Function Calls Inspected by memstomp

| Function | Description |
|-----------------|---|
| memcpy | Copies n bytes from one memory area to another and returns a pointer to the second memory area. |
| memccpy | Copies a maximum of n bytes from one memory area to another and stops when a certain character is found. It either returns a pointer to the byte following the last written byte, or NULL if the given character is not found. |
| mempcpy | Copies n bytes from one memory area to another and returns a pointer to the byte following the last written byte. |
| strcpy | Copies a string from one memory area to another and returns a pointer to the second string. |
| stpcpy | Copies a string from one memory area to another and returns a pointer to the terminating null byte of the second string. |
| strncpy | Copies a maximum of n characters from one string to another and returns a pointer to the second string. |
| stpncpy | Copies a maximum of n characters from one string to another. It either returns a pointer to the terminating null byte of the second string, or if the string is not null-terminated, a pointer to the byte following the last written byte. |
| strcat | Appends one string to another while overwriting the terminating null byte of the second string and adding a new one at its end. It returns a pointer to the new string. |
| strncat | Appends a maximum of n characters from one string to another while overwriting the terminating null byte of the second string and adding a new one at its end. It returns a pointer to the new string. |
| wmemcpy | The wide-character equivalent of the memcpy() function that copies n wide characters from one array to another and returns a pointer to the second array. |
| wmempcpy | The wide-character equivalent of the mempcpy() function that copies n wide characters from one array to another and returns a pointer to the byte following the last written wide character. |
| wscpy | The wide-character equivalent of the strcpy() function that copies a wide-character string from one array to another and returns a pointer to the second array. |
| wcsncpy | The wide-character equivalent of the strncpy() function that copies a maximum of n wide characters from one array to another and returns a pointer to the second string. |
| wscat | The wide-character equivalent of the strcat() function that appends one wide-character string to another while overwriting the terminating null byte of the second string and adding a new one at its end. It returns a pointer to the new string. |
| wcsncat | The wide-character equivalent of the strncat() function that appends a maximum of n wide characters from one array to another while overwriting the terminating null byte of the second wide-character string and adding a new one at its end. It returns a pointer to the new string. |

10.1. Installing memstomp

In Red Hat Developer Toolset, the **memstomp** utility is provided by the *devtoolset-2-memstomp* package and is automatically installed with *devtoolset-2-toolchain* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).



Note

The Red Hat Developer Toolset version of **memstomp** is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures.

10.2. Using memstomp

To run the **memstomp** utility on a program you want to analyze, type the following at a shell prompt:

```
sc1 enable devtoolset-2 'memstomp program [argument...]
```

To immediately terminate the analyzed program when a problem is detected, run the utility with the **--kill** (or **-k** for short) command line option:

```
sc1 enable devtoolset-2 'memstomp --kill program [argument...]
```

The use of the **--kill** option is especially recommended if you are analyzing a multi-threaded program; the internal implementation of backtraces is not thread-safe and running the **memstomp** utility on a multi-threaded program without this command line option can therefore produce unreliable results.

Additionally, if you have compiled the analyzed program with the debugging information or this debugging information is available to you, you can use the **--debug-info** (or **-d**) command line option to produce a more detailed backtrace:

```
sc1 enable devtoolset-2 'memstomp --debug-info program [argument...]
```

For detailed instructions on how to compile your program with the debugging information built in the binary file, see [Section 8.2, “Preparing a Program for Debugging”](#). For information on how to install debugging information for any of the Red Hat Developer Toolset packages, see [Section 1.5.4, “Installing Debugging Information”](#).

Note that you can execute any command using the **sc1** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **memstomp** as default:

```
sc1 enable devtoolset-2 'bash'
```

Example 10.1. Using memstomp

In the current working directory, create a source file named **employee.c** with the following contents:

```
#include <stdio.h>
#include <string.h>

#define BUFSIZE 80

int main(int argc, char *argv[]) {
    char employee[BUFSIZE] = "John,Doe,john@example.com";
    char name[BUFSIZE] = {0};
    char surname[BUFSIZE] = {0};
    char *email;
    size_t length;

    /* Extract the information: */
    memccpy(name, employee, ',', BUFSIZE);
    length = strlen(name);
    memccpy(surname, employee + length, ',', BUFSIZE);
    length += strlen(surname);
    email = employee + length;

    /* Compose the new entry: */
    strcat(employee, surname);
    strcpy(employee, name);
    strcat(employee, email);

    /* Print the result: */
    puts(employee);

    return 0;
}
```

Compile this program into a binary file named **employee** by using the following command:

```
~]$ scl enable devtoolset-2 'gcc -rdynamic -g -o employee employee.c'
```

To identify erroneous function calls with overlapping memory regions, type:

```
~]$ scl enable devtoolset-2 'memstomp --debug-info ./employee'
memstomp: 0.1.4 successfully initialized for process employee (pid 14887).
```

```
strcat(dest=0x7fff13afc265, src=0x7fff13afc269, bytes=21) overlap for
employee(14887)
```

```
??:0    strcpy()
??:0    strcpy()
??:0    _Exit()
??:0    strcat()
employee.c:26  main()
??:0    __libc_start_main()
??:0    _start()
```

```
John,john@example.comDoe,
```

10.3. Additional Resources

A detailed description of **memstomp** and its features is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

- **memstomp(1)** — The manual page for the **memstomp** utility provides detailed information about its usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man memstomp'
```

See Also

- [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 8, GNU Debugger \(GDB\)](#) provides information on how to debug programs written in C, C++, and Fortran.
- [Chapter 9, strace](#) documents how to use the *strace* utility to monitor system calls that a program uses and signals it receives.
- [Chapter 12, Valgrind](#) explains how to use *valgrind* to profile applications and detect memory errors and memory management problems, such as the use of uninitialized memory, improper allocation and freeing of memory, and the use of improper arguments in system calls.

Part VI. Performance Monitoring Tools

Chapter 11. SystemTap

SystemTap is a tracing and probing tool that allows users to monitor the activities of the entire system without needing to instrument, recompile, install, and reboot. It is programmable with a custom scripting language, which gives it expressiveness (to trace, filter, and analyze) and reach (to look into the running kernel and applications).

SystemTap can monitor various types of events, such as function calls within the kernel or applications, timers, tracepoints, performance counters, and so on. Some included example scripts produce output similar to **netstat**, **ps**, **top**, and **iostat**, others include pretty-printed function callgraph traces or tools for working around security bugs.

Red Hat Developer Toolset is distributed with **SystemTap 2.1**.

Table 11.1. Tools Distributed with SystemTap for Red Hat Developer Toolset

| Name | Description |
|--------------------|--|
| stap | Translates probing instructions into C code, builds a kernel module, and loads it into a running Linux kernel. |
| staprun | Loads, unloads, attaches to, and detaches from kernel modules built with the stap utility. |
| stapsh | Serves as a remote shell for SystemTap. |
| stap-prep | Determines and—if possible—downloads the kernel information packages that are required to run SystemTap. |
| stap-merge | Merges per-CPU files. This script is automatically executed when the stap utility is executed with the -b command line option. |
| stap-report | Gathers important information about the system for the purpose of reporting a bug in SystemTap. |

11.1. Installing SystemTap

In Red Hat Developer Toolset, **SystemTap** is provided by the *devtoolset-2-systemtap* package and is automatically installed with *devtoolset-2-perftools* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).



Note

The Red Hat Developer Toolset version of SystemTap is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures.

In order to place instrumentation into the Linux kernel, SystemTap may also require installation of additional packages with debugging information. To determine which packages to install, run the **stap-prep** utility as follows:

```
scl enable devtoolset-2 'stap-prep'
```

Note that if you execute this command as the **root** user, the utility automatically offers the packages for installation. For more information on how to install these packages on your system, see the *Red Hat Enterprise Linux SystemTap Beginners Guide*.

11.2. Using SystemTap

To execute any of the tools that are part of SystemTap, type the following at a shell prompt:

```
scl enable devtoolset-2 'tool [option...]
```

See [Table 11.1, “Tools Distributed with SystemTap for Red Hat Developer Toolset”](#) for a list of tools that are distributed with SystemTap. For example, to run the **stap** tool to build an instrumentation module, type:

```
scl enable devtoolset-2 'stap [option...] argument...'
```

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset SystemTap as default:

```
scl enable devtoolset-2 'bash'
```



Note

To verify the version of SystemTap you are using at any point, type the following at a shell prompt:

```
which stap
```

Red Hat Developer Toolset's **stap** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset SystemTap:

```
stap -V
```

11.3. Additional Resources

A detailed description of SystemTap and its features is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

- **stap(1)** — The manual page for the **stap** command provides detailed information on its usage, as well as references to other related manual pages. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man stap'
```

- **staprun(8)** — The manual page for the **staprun** command provides detailed information on its usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man staprun'
```

- » *SystemTap Tapset Reference Manual* — HTML documentation on the most common tapset definitions is located at `/opt/rh/devtoolset-2/root/usr/share/doc/devtoolset-2-systemtap-client-2.1/index.html`.

Online Documentation

- » [Red Hat Enterprise Linux 6 SystemTap Beginners Guide](#) — The *SystemTap Beginners Guide* for Red Hat Enterprise Linux 6 provides an introduction to SystemTap and its usage.
- » [Red Hat Enterprise Linux 5 SystemTap Beginners Guide](#) — The *SystemTap Beginners Guide* for Red Hat Enterprise Linux 5 provides an introduction to SystemTap and its usage.
- » [Red Hat Enterprise Linux 6 SystemTap Tapset Reference](#) — The *SystemTap Tapset Reference* for Red Hat Enterprise Linux 6 provides further details about SystemTap.
- » [Red Hat Enterprise Linux 5 SystemTap Tapset Reference](#) — The *SystemTap Tapset Reference* for Red Hat Enterprise Linux 5 provides further details about SystemTap.
- » [The SystemTap Documentation](#) — The official SystemTap documentation provides further documentation on SystemTap, as well as numerous examples of SystemTap scripts.

See Also

- » [Section B.8, “Changes in SystemTap”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of SystemTap and the version distributed in the previous release of Red Hat Developer Toolset.
- » [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- » [Chapter 12, Valgrind](#) explains how to use Valgrind to profile applications and detect memory errors and memory management problems, such as the use of uninitialized memory, improper allocation and freeing of memory, and the use of improper arguments in system calls.
- » [Chapter 13, OProfile](#) explains how to use OProfile to determine which sections of code consume the greatest amount of CPU time and why.
- » [Chapter 14, Dyninst](#) documents how to use the Dyninst library to instrument a user-space executable.

Chapter 12. Valgrind

Valgrind is an instrumentation framework that ships with a number of tools to profile applications. It can be used to detect various memory errors and memory management problems, such as the use of uninitialized memory or improper allocation and freeing of memory, or to identify the use of improper arguments in system calls. For a complete list of profiling tools that are distributed with the Red Hat Developer Toolset version of Valgrind, see [Table 12.1, “Tools Distributed with Valgrind for Red Hat Developer Toolset”](#).

Valgrind profiles an application by rewriting it and instrumenting the rewritten binary. This allows you to profile your application without the need to recompile it, but it also makes Valgrind significantly slower than other profilers, especially when performing extremely detailed runs. It is therefore not suited to debugging time-specific issues, or kernel-space debugging.

Red Hat Developer Toolset is distributed with **Valgrind 3.8.1**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements.

Table 12.1. Tools Distributed with Valgrind for Red Hat Developer Toolset

| Name | Description |
|-------------------|---|
| Memcheck | Detects memory management problems by intercepting system calls and checking all read and write operations. |
| Cachegrind | Identifies the sources of cache misses by simulating the level 1 instruction cache (I1), level 1 data cache (D1), and unified level 2 cache (L2). |
| Callgrind | Generates a call graph representing the function call history. |
| Helgrind | Detects synchronization errors in multithreaded C, C++, and Fortran programs that use POSIX threading primitives. |
| DRD | Detects errors in multithreaded C and C++ programs that use POSIX threading primitives or any other threading concepts that are built on top of these POSIX threading primitives. |
| Massif | Monitors heap and stack usage. |

12.1. Installing Valgrind

In Red Hat Developer Toolset, Valgrind is provided by the *devtoolset-2-valgrind* package and is automatically installed with *devtoolset-2-perftools*. If you intend to use Valgrind to profile parallel programs that use the Message Passing Interface (MPI) protocol, also install the *devtoolset-2-valgrind-openmpi* package by typing the following at a shell prompt as **root**:

```
yum install devtoolset-2-valgrind-openmpi
```

For detailed instructions on how to install Red Hat Developer Toolset and related packages to your system, see [Section 1.5, “Installing Red Hat Developer Toolset”](#).

**Note**

The Red Hat Developer Toolset version of Valgrind is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures. Note that if you use Valgrind in combination with the GNU Debugger, it is recommended that you use the version of GDB that is included in Red Hat Developer Toolset to ensure that all features are fully supported.

12.2. Using Valgrind

To run any of the Valgrind tools on a program you want to profile, type the following at a shell prompt:

```
sc1 enable devtoolset-2 'valgrind [--tool=tool] program [argument...]
```

See [Table 12.1, “Tools Distributed with Valgrind for Red Hat Developer Toolset”](#) for a list of tools that are distributed with Valgrind. The argument of the **--tool** command line option must be specified in lower case, and if this option is omitted, Valgrind uses **Memcheck** by default. For example, to run Cachegrind on a program to identify the sources of cache misses, type:

```
sc1 enable devtoolset-2 'valgrind --tool=cachegrind program [argument...]
```

Note that you can execute any command using the **sc1** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset Valgrind as default:

```
sc1 enable devtoolset-2 'bash'
```

**Note**

To verify the version of Valgrind you are using at any point, type the following at a shell prompt:

```
which valgrind
```

Red Hat Developer Toolset's **valgrind** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset Valgrind:

```
valgrind --version
```

12.3. Additional Resources

A detailed description of Valgrind and its features is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

- **valgrind(1)** — The manual page for the **valgrind** utility provides detailed information on how to

use Valgrind. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man valgrind'
```

- *Valgrind Documentation* — HTML documentation for Valgrind is located at `/opt/rh/devtoolset-2/root/usr/share/doc/devtoolset-2-valgrind-3.8.1/html/index.html`.

Online Documentation

- [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information about Valgrind and its Eclipse plug-in.
- [Red Hat Enterprise Linux 6 Performance Tuning Guide](#) — The *Performance Tuning Guide* for Red Hat Enterprise Linux 6 provides more detailed information about using Valgrind to profile applications.

See Also

- [Section A.6, “Changes in Valgrind”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of Valgrind and the version distributed in the previous release of Red Hat Developer Toolset.
- [Chapter 1, *Red Hat Developer Toolset*](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 10, *memstomp*](#) documents how to use the memstomp utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.
- [Chapter 11, *SystemTap*](#) provides an introduction to SystemTap and explains how to use it to monitor the activities of a running system.
- [Chapter 13, *OProfile*](#) explains how to use OProfile to determine which sections of code consume the greatest amount of CPU time and why.
- [Chapter 14, *Dyninst*](#) documents how to use the Dyninst library to instrument a user-space executable.

Chapter 13. OProfile

OProfile is a low overhead, system-wide profiler that uses the performance monitoring hardware on the processor to retrieve information about the kernel and executables on the system, such as when memory is referenced, the number of level 2 cache (L2) requests, and the number of hardware interrupts received. It consists of a configuration utility, a daemon for collecting data, and a number of tools that can be used to transform the data into a human-readable form. For a complete list of tools that are distributed with the Red Hat Developer Toolset version of OProfile, see [Table 13.1, “Tools Distributed with OProfile for Red Hat Developer Toolset”](#).

OProfile profiles an application without adding any instrumentation by recording the details of every nth event. This allows it to consume fewer resources than Valgrind, but also causes its samples to be less precise. Unlike Valgrind, which only collects data for a single process and its children in user-space, OProfile is well suited to collect system-wide data on both user-space and kernel-space processes, and requires **root** privileges to run.

Red Hat Developer Toolset is distributed with **OProfile 0.9.8**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements.

Table 13.1. Tools Distributed with OProfile for Red Hat Developer Toolset

| Name | Description |
|-------------------|---|
| oprofiled | The OProfile daemon that collects profiling data. |
| opcontrol | Starts, stops, and configures the OProfile daemon. |
| opannotate | Generates an annotated source file or assembly listing from the profiling data. |
| oparchive | Generates a directory containing executable, debug, and sample files. |
| opgprof | Generates a summary of a profiling session in a format compatible with gprof . |
| ophelp | Displays a list of available events. |
| opimport | Converts a sample database file from a foreign binary format to the native format. |
| opjitconv | Converts a just-in-time (JIT) dump file to the Executable and Linkable Format (ELF). |
| opreport | Generates image and symbol summaries of a profiling session. |

13.1. Installing OProfile

In Red Hat Developer Toolset, OProfile is provided by the *devtoolset-2-oprofile* package and is automatically installed with *devtoolset-2-perftools* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).



Note

The Red Hat Developer Toolset version of OProfile is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures. Note that the **operf** utility relies on a kernel feature that is currently not available in Red Hat Enterprise Linux 5, and as a consequence, this utility only works and is supported on Red Hat Enterprise Linux 6.

13.2. Using OProfile

To run any of the tools that are distributed with OProfile, type the following at a shell prompt as **root**:

```
scl enable devtoolset-2 'tool [option...]
```

See [Table 13.1, “Tools Distributed with OProfile for Red Hat Developer Toolset”](#) for a list of tools that are distributed with OProfile. For example, to use the **ophelp** command to list available events in the XML format, type:

```
scl enable devtoolset-2 'ophelp -X'
```

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset OProfile as default:

```
scl enable devtoolset-2 'bash'
```



Note

To verify the version of OProfile you are using at any point, type the following at a shell prompt:

```
which opcontrol
```

Red Hat Developer Toolset's **opcontrol** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset OProfile:

```
opcontrol --version
```

13.3. Additional Resources

A detailed description of OProfile and its features is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

- **oprofile(1)** — The manual page named **oprofile** provides an overview of OProfile and available tools. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man oprofile'
```

- **opannotate(1)**, **oparchive(1)**, **opcontrol(1)**, **opggprof(1)**, **ophelp(1)**, **opimport(1)**, **opreport(1)** — Manual pages for various tools distributed with OProfile provide more information on their respective usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-2 'man tool'
```

Online Documentation

- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information on OProfile.
- ▶ [Red Hat Enterprise Linux 6 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 6 describes in detail how to install, configure, and start using OProfile on this system.
- ▶ [Red Hat Enterprise Linux 5 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 5 describes in detail how to install, configure, and start using OProfile on this system.

See Also

- ▶ [Section A.7, “Changes in OProfile”](#) provides a comprehensive list of changes and improvements over the Red Hat Enterprise Linux system version of OProfile and the version distributed in the previous release of Red Hat Developer Toolset.
- ▶ [Chapter 1, *Red Hat Developer Toolset*](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 11, *SystemTap*](#) provides an introduction to SystemTap and explains how to use it to monitor the activities of a running system.
- ▶ [Chapter 12, *Valgrind*](#) explains how to use Valgrind to profile applications and detect memory errors and memory management problems, such as the use of uninitialized memory, improper allocation and freeing of memory, and the use of improper arguments in system calls.
- ▶ [Chapter 14, *Dyninst*](#) documents how to use the Dyninst library to instrument a user-space executable.

Chapter 14. Dyninst

The **Dyninst** library provides an application programming interface (API) for instrumenting and working with user-space executables during their execution. It can be used to insert code into a running program, change certain subroutine calls, or even remove them from the program, and serves as a valuable debugging and performance monitoring tool. The Dyninst API is also commonly used along with SystemTap to allow non-root users to instrument user-space executables.

Red Hat Developer Toolset is distributed with **Dyninst 8.0**.

14.1. Installing Dyninst

In Red Hat Developer Toolset, the **Dyninst** library is provided by the *devtoolset-2-dyninst* package and is automatically installed with *devtoolset-2-perftools* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#). In addition, it is recommended that you also install the GNU Compiler Collection provided by the *devtoolset-2-toolchain* package.



Note

The Red Hat Developer Toolset version of Dyninst is available for both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures.

If you intend to write a custom instrumentation for binaries, install the relevant header files by running the following command as **root**:

```
yum install devtoolset-2-dyninst-devel
```

You can also install API documentation for this library by typing the following at a shell prompt as **root**:

```
yum install devtoolset-2-dyninst-doc
```

For a complete list of documents that are included in the *devtoolset-2-dyninst-doc* package, see [Section 14.3, “Additional Resources”](#). For detailed instructions on how to install optional packages to your system, see [Section 1.5, “Installing Red Hat Developer Toolset”](#).

14.2. Using Dyninst

14.2.1. Using Dyninst with SystemTap

To use Dyninst along with SystemTap to allow non-root users to instrument user-space executables, run the **stap** command with the **--dyninst** (or **--runtime=dyninst**) command line option. This tells **stap** to translate a SystemTap script into C code that uses the Dyninst library, compile this C code into a shared library, and then load the shared library and run the script. Note that when executed like this, the **stap** command also requires the **-c** or **-x** command line option to be specified.

To use the Dyninst runtime to instrument an executable file, type the following at a shell prompt:

```
scl enable devtoolset-2 "stap --dyninst -c 'command' [option...]
[argument...]"
```

Similarly, to use the Dyninst runtime to instrument a user's process, type:

```
scl enable devtoolset-2 "stap --dyninst -x process_id [option...]  
[argument...]"
```

See [Chapter 11, SystemTap](#) for more information about the Red Hat Developer Toolset version of SystemTap. For a general introduction to SystemTap and its usage, see the *SystemTap Beginners Guide* for Red Hat Enterprise Linux.

Example 14.1. Using Dyninst with SystemTap

Consider a source file named **exercise.C** that has the following contents:

```
#include <stdio.h>

void print_iteration(int value) {
    printf("Iteration number %d\n", value);
}

int main(int argc, char **argv) {
    int i;
    printf("Enter the starting number: ");
    scanf("%d", &i);
    for(; i>0; --i)
        print_iteration(i);
    return 0;
}
```

This program prompts the user to enter a starting number and then counts down to 1, calling the **print_iteration()** function for each iteration in order to print the number to standard output. To compile this program on the command line using the **g++** compiler from Red Hat Developer Toolset, type the following at a shell prompt:

```
~]$ scl enable devtoolset-2 'g++ -g -o exercise exercise.C'
```

Now consider another source file named **count.stp** with the following contents:

```
#!/usr/bin/stap

global count = 0

probe process.function("print_iteration") {
    count++
}

probe end {
    printf("Function executed %d times.\n", count)
}
```

This SystemTap script prints the total number of times the **print_iteration()** function was called during the execution of a process. To run this script on the **exercise** binary file, type:

```
~]$ scl enable devtoolset-2 "stap --dyninst -c './exercise' count.stp"
Enter the starting number: 5
Iteration number 5
Iteration number 4
Iteration number 3
Iteration number 2
Iteration number 1
Function executed 5 times.
```

14.2.2. Using Dyninst as a Stand-alone Application

Before using the Dyninst library as a stand-alone application, set the value of the **DYNINSTAPI_RT_LIB** environment variable to the path to the runtime library file. If you are running a 64-bit system, you can do so by typing the following at a shell prompt:

```
export DYNINSTAPI_RT_LIB=/opt/rh/devtoolset-  
2/root/usr/lib64/dyninst/libdyninstAPI_RT.so.8.0
```

If you are running a 32-bit system, type:

```
export DYNINSTAPI_RT_LIB=/opt/rh/devtoolset-  
2/root/usr/lib/dyninst/libdyninstAPI_RT.so.8.0
```

This sets the **DYNINSTAPI_RT_LIB** environment variable in the current shell session.

[Example 14.2, “Using Dyninst as a Stand-alone Application”](#) illustrates how to write and build a program to monitor the execution of a user-space process. For a detailed explanation of how to use Dyninst, see the resources listed in [Section 14.3, “Additional Resources”](#).

Example 14.2. Using Dyninst as a Stand-alone Application

Consider the **exercise.C** source file from [Example 14.1, “Using Dyninst with SystemTap”](#): this program prompts the user to enter a starting number and then counts down to 1, calling the **print_iteration()** function for each iteration in order to print the number to standard output.

Now consider another source file named **count.C** with the following contents:

```

#include <stdio.h>
#include <fcntl.h>
#include "BPatch.h"
#include "BPatch_process.h"
#include "BPatch_function.h"
#include "BPatch_Vector.h"
#include "BPatch_thread.h"
#include "BPatch_point.h"

void usage() {
    fprintf(stderr, "Usage: count <process_id> <function>\n");
}

// Global information for counter
BPatch_variableExpr *counter = NULL;

void createCounter(BPatch_process *app, BPatch_image *appImage) {
    int zero = 0;
    counter = app->malloc(*appImage->findType("int"));
    counter->writeValue(&zero);
}

bool interceptfunc(BPatch_process *app,
                  BPatch_image *appImage,
                  char *funcName) {
    BPatch_Vector<BPatch_function *> func;
    appImage->findFunction(funcName, func);
    if(func.size() == 0) {
        fprintf(stderr, "Unable to find function to instrument()\n");
        exit (-1);
    }
    BPatch_Vector<BPatch_snippet *> incCount;
    BPatch_Vector<BPatch_point *> *points;
    points = func[0]->findPoint(BPatch_entry);
    if ((*points).size() == 0) {
        exit (-1);
    }

    BPatch_arithExpr counterPlusOne(BPatch_plus, *counter, BPatch_constExpr(1));
    BPatch_arithExpr addCounter(BPatch_assign, *counter, counterPlusOne);

    return app->insertSnippet(addCounter, *points);
}

void printCount(BPatch_thread *thread, BPatch_exitType) {
    int val = 0;
    counter->readValue(&val, sizeof(int));
    fprintf(stderr, "Function executed %d times.\n", val);
}

BPatch bpatch;

int main(int argc, char *argv[]) {
    int pid;
    if (argc != 3) {
        usage();
        exit(1);
    }
    pid = atoi(argv[1]);

```

```

BPatch_process *app = bpatch.processAttach(NULL, pid);
if (!app) exit (-1);
BPatch_image *appImage = app->getImage();
createCounter(app, appImage);
fprintf(stderr, "Finding function %s(): ", argv[2]);
BPatch_Vector<BPatch_function*> countFuncs;
fprintf(stderr, "OK\nInstrumenting function %s(): ", argv[2]);
interceptfunc(app, appImage, argv[2]);
bpatch.registerExitCallback(printCount);
fprintf(stderr, "OK\nWaiting for process %d to exit...\n", pid);
app->continueExecution();
while (!app->isTerminated())
    bpatch.waitForStatusChange();
return 0;
}

```

This program accepts a process ID and a function name as command line arguments and then prints the total number of times the function was called during the execution of the process. You can use the following **Makefile** to build these two files:

```

DTS      = /opt/rh/devtoolset-2/root
CXXFLAGS = -g -I$(DTS)/usr/include/dyninst
LBITS    := $(shell getconf LONG_BIT)

ifeq ($(LBITS),64)
    DYNINSTLIBS = $(DTS)/usr/lib64/dyninst
else
    DYNINSTLIBS = $(DTS)/usr/lib/dyninst
endif

.PHONY: all
all: count exercise

count: count.C
g++ $(CXXFLAGS) count.C -I /usr/include/dyninst -c
g++ $(CXXFLAGS) count.o -L $(DYNINSTLIBS) -ldyninstAPI -o count

exercise: exercise.C
g++ $(CXXFLAGS) exercise.C -o exercise

.PHONY: clean
clean:
rm -rf *~ *.o count exercise

```

To compile the two programs on the command line using the **g++** compiler from Red Hat Developer Toolset, run the **make** utility as follows:

```

~]$ scl enable devtoolset-2 make
g++ -g -I/opt/rh/devtoolset-2/root/usr/include/dyninst count.C -c
g++ -g -I/opt/rh/devtoolset-2/root/usr/include/dyninst count.o -L
/opt/rh/devtoolset-2/root/usr/lib64/dyninst -ldyninstAPI -o count
g++ -g -I/opt/rh/devtoolset-2/root/usr/include/dyninst exercise.C -o exercise

```

This creates new binary files called **exercise** and **count** in the current working directory.

In one shell session, execute the **exercise** binary file as follows and wait for it to prompt you to enter the starting number:

```
~]$ ./exercise
Enter the starting number:
```

Do not enter this number. Instead, start another shell session and type the following at its prompt to set the **DYNINSTAPI_RT_LIB** environment variable and execute the **count** binary file:

```
~]$ export DYNINSTAPI_RT_LIB=/opt/rh/devtoolset-2/root/usr/lib64/dyninst/libdyninstAPI_RT.so.8.0
~]$ ./count `pidof exercise` print_iteration
Finding function print_iteration(): OK
Instrumenting function print_iteration(): OK
Waiting for process 8607 to exit...
```

Now switch back to the first shell session and enter the starting number as requested by the **exercise** program. For example:

```
Enter the starting number: 5
Iteration number 5
Iteration number 4
Iteration number 3
Iteration number 2
Iteration number 1
```

When the **exercise** program terminates, the **count** program displays the number of times the **print_iteration()** function was executed:

```
Function executed 5 times.
```

14.3. Additional Resources

A detailed description of Dyninst and its features is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

The *devtoolset-2-dyninst-doc* package installs the following documents in the **/opt/rh/devtoolset-2/root/usr/share/doc/devtoolset-2-dyninst-doc-8.0/** directory:

- » *Dyninst Programmer's Guide* — A detailed description of the Dyninst API is stored in the **DyninstAPI.pdf** file.
- » *DynC API Programmer's Guide* — An introduction to DynC API is stored in the **dynC_API.pdf** file.
- » *ParseAPI Programmer's Guide* — An introduction to the ParseAPI is stored in the **ParseAPI.pdf** file.
- » *PatchAPI Programmer's Guide* — An introduction to PatchAPI is stored in the **PatchAPI.pdf** file.
- » *ProcControlAPI Programmer's Guide* — A detailed description of ProcControlAPI is stored in the **ProcControlAPI.pdf** file.
- » *StackwalkerAPI Programmer's Guide* — A detailed description of StackwalkerAPI is stored in the **stackwalker.pdf** file.
- » *SymtabAPI Programmer's Guide* — An introduction to SymtabAPI is stored in the **SymtabAPI.pdf** file.

- *InstructionAPI Reference Manual* — A detailed description of the InstructionAPI is stored in the **InstructionAPI.pdf** file.

For information on how to install this package on your system, see [Section 14.1, “Installing Dyninst”](#).

Online Documentation

- [Dyninst Home Page](#) — The project home page provides links to additional documentation and related publications.
- [Red Hat Enterprise Linux 6 SystemTap Beginners Guide](#) — The *SystemTap Beginners Guide* for Red Hat Enterprise Linux 6 provides an introduction to SystemTap and its usage.
- [Red Hat Enterprise Linux 5 SystemTap Beginners Guide](#) — The *SystemTap Beginners Guide* for Red Hat Enterprise Linux 5 provides an introduction to SystemTap and its usage.
- [Red Hat Enterprise Linux 6 SystemTap Tapset Reference](#) — The *SystemTap Tapset Reference* for Red Hat Enterprise Linux 6 provides further details about SystemTap.
- [Red Hat Enterprise Linux 5 SystemTap Tapset Reference](#) — The *SystemTap Tapset Reference* for Red Hat Enterprise Linux 5 provides further details about SystemTap.

See Also

- [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 11, SystemTap](#) provides an introduction to SystemTap and explains how to use it to monitor the activities of a running system.
- [Chapter 12, Valgrind](#) explains how to use Valgrind to profile applications and detect memory errors and memory management problems, such as the use of uninitialized memory, improper allocation and freeing of memory, and the use of improper arguments in system calls.
- [Chapter 13, OProfile](#) explains how to use OProfile to determine which sections of code consume the greatest amount of CPU time and why.

Part VII. Getting Help

Chapter 15. Accessing Red Hat Product Documentation

Red Hat Product Documentation located at <https://access.redhat.com/site/documentation/> serves as a central source of information. It is currently translated in 22 languages and for each product, it provides different kinds of books from release and technical notes to installation, user, and reference guides in HTML, PDF, and EPUB formats.

Below is a brief list of documents that are directly or indirectly relevant to this book.

Red Hat Developer Toolset

- ▶ [Red Hat Developer Toolset 2.1 Release Notes](#) — The *Release Notes* for Red Hat Developer Toolset 2.1 provide more information about this product.
- ▶ [Red Hat Developer Toolset 2.1 Software Collections Guide](#) — The *Software Collections Guide* for Red Hat Developer Toolset 2.1 explains the concept of Software Collections and documents the `sc1` tool.

Red Hat Enterprise Linux

- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides detailed information about libraries and runtime support, compiling and building, debugging, and profiling.
- ▶ [Red Hat Enterprise Linux 6 Installation Guide](#) — The *Installation Guide* for Red Hat Enterprise Linux 6 explains how to obtain, install, and update the system.
- ▶ [Red Hat Enterprise Linux 5 Installation Guide](#) — The *Installation Guide* for Red Hat Enterprise Linux 5 explains how to obtain, install, and update the system.
- ▶ [Red Hat Enterprise Linux 6 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 6 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 6.
- ▶ [Red Hat Enterprise Linux 5 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 5 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 5.

Chapter 16. Contacting Global Support Services

Unless you have a Self-Support subscription, when both the Red Hat Documentation website and Customer Portal fail to provide the answers to your questions, you can contact **Global Support Services (GSS)**.

16.1. Gathering Required Information

Several items of information should be gathered before contacting GSS.

Background Information

Ensure you have the following background information at hand before calling GSS:

- Hardware type, make, and model on which the product runs
- Software version
- Latest upgrades
- Any recent changes to the system
- An explanation of the problem and the symptoms
- Any messages or significant information about the issue



Note

If you ever forget your Red Hat login information, it can be recovered at <https://access.redhat.com/site/help/LoginAssistance.html>.

Diagnostics

The diagnostics report for Red Hat Enterprise Linux is required as well. This report is also known as a *sosreport* and the program to create the report is provided by the *sos* package. To install the *sos* package and all its dependencies on your system, type the following at a shell prompt as **root**:

```
yum install sos
```

To generate the report, run as **root**:

```
sosreport
```

For more information, access the Knowledgebase article at <https://access.redhat.com/kb/docs/DOC-3593>.

Account and Contact Information

In order to help you, GSS requires your account information to customize their support, as well contact information to get back to you. When you contact GSS ensure you have your:

- Red Hat customer number or Red Hat Network (RHN) login name
- Company name
- Contact name
- Preferred method of contact (phone or email) and contact information (phone number or email address)

Issue Severity

Determining an issue's severity is important to allow the GSS team to prioritize their work. There are four levels of severity.

Severity 1 (urgent)

A problem that severely impacts your use of the software for production purposes. It halts your business operations and has no procedural work around.

Severity 2 (high)

A problem where the software is functioning but production is severely reduced. It causes a high impact to business operations and no work around exists.

Severity 3 (medium)

A problem that involves partial, non-critical loss of the use of the software. There is a medium to low impact on your business and business continues to function by utilizing a work around.

Severity 4 (low)

A general usage question, report of a documentation error, or a recommendation for a future product improvement.

For more information on determining the severity level of an issue, see

<https://access.redhat.com/support/policy/severity>.

Once the issue severity has been determined, submit a service request through the Customer Portal under the **Connect** option, or at <https://access.redhat.com/support/contact/technicalSupport.html>. Note that you need your Red Hat login details in order to submit service requests.

If the severity is level 1 or 2, then follow up your service request with a phone call. Contact information and business hours are found at <https://access.redhat.com/support/contact/technicalSupport.html>.

If you have a premium subscription, then after hours support is available for Severity 1 and 2 cases.

Turn-around rates for both premium subscriptions and standard subscription can be found at <https://access.redhat.com/support/offerings/production/sla.html>.

16.2. Escalating an Issue

If you feel as though an issue is not being handled correctly or adequately, you can escalate it. There are two types of escalations:

Technical escalation

If an issue is not being resolved appropriately or if you need a more senior resource to attend to it.

Management escalation

If the issue has become more severe or you believe it requires a higher priority.

More information on escalation, including contacts, is available at

https://access.redhat.com/support/policy/mgt_escalation.html.

16.3. Re-opening a Service Request

If more relevant information regarding a closed service request (such as the problem reoccurring), you can re-open it via the Red Hat Customer Portal at

https://access.redhat.com/support/policy/mgt_escalation.html or by calling your local support center, the details of which can be found at <https://access.redhat.com/support/contact/technicalSupport.html>.



Important

In order to re-open a service request, you need the original service request number.

16.4. Additional Resources

For more information, see the resources listed below.

Online Documentation

- ▶ [Getting Started](#) — The *Getting Started* page serves as a starting point for people who purchased a Red Hat subscription, and offers the *Red Hat Welcome Kit* and *Quick Guide to Red Hat Support* for download.
- ▶ [How can a RHEL Self-Support subscription be used?](#) — A Knowledgebase article for customers with a Self-Support subscription.
- ▶ [Red Hat Global Support Services and public mailing lists](#) — A Knowledgebase article that answers frequent questions about public Red Hat mailing lists.

Changes in Version 2.1

The sections below document features and compatibility changes introduced in Red Hat Developer Toolset 2.1.

A.1. Changes in Eclipse

Red Hat Developer Toolset 2.1 is distributed with **Eclipse 4.3.1** and other plug-ins from the Eclipse Foundation's 2013 Eclipse Kepler 4.3 Service Release 1. This version is newer than the version of Eclipse distributed with the previous release of Red Hat Developer Toolset, and provides a number of bug fixes and feature enhancements.

Below is a comprehensive list of new features and compatibility changes in this release. For details on how to use these new features, see the built-in Eclipse documentation.

A.1.1. Changes Since Red Hat Developer Toolset 2.0

The following features have been added since the release of the Eclipse Platform and native plug-ins included in Red Hat Developer Toolset 2.0:

- ▶ The **Eclipse Platform** has been updated from version 4.3.0 to 4.3.1. This service release provides over 100 bug fixes and optimizations.
- ▶ The **Eclipse C/C++ Development Toolkit** (CDT) has been updated from version 8.2.0 to 8.2.1. This service version update provides over 30 bug fixes and optimizations, including fixes related to the organization of **include** statements.
- ▶ **Mylyn**, a set of task and application life cycle management plug-ins, has been updated from version 3.9 to 3.9.1. This update provides over 40 bug fixes and optimizations, including various improvements to the Bugzilla connector.
- ▶ **EGit**, a Git integration plug-in for Eclipse, has been updated from version 2.3.1 to 3.1. This new major release includes:
 - support for Mylin links in the **History** view, **Git Staging** view, and **Commit** dialog window;
 - improved Gerrit integration;
 - a significantly enhanced **Git Staging** view;
 - an enhanced command line tool (**jgit**);
 - several performance improvements.
- ▶ The **Eclipse Linux Tools** plug-in collection has been updated from version 2.0 to 2.2. This update includes:
 - support for the **perf** command line tool in the **OProfile** plug-in; the **perf** command is now used by default to make generic setup easier;
 - an enhanced **RPM Stubby** plug-in, which is now able to generate RPM spec files from Perl makefiles, Ruby gamespec files, and Python setup.py files;
 - an enhanced **Spec File Editor** plug-in, which is now able to download source files, prepare a buildroot, and also provides a **Build** menu, a corresponding toolbar, and improved support for hyperlinks in the editor;
 - an enhanced **SystemTap .stp** files editor, which now has improved formatting capabilities;
 - improved **SystemTap** graphing, which has been streamlined to be easier to launch and configure.

A.2. Changes in GCC

Red Hat Developer Toolset 2.1 is distributed with **GCC 4.8.2**, which provides a number of bug fixes over version 4.8.1 included in Red Hat Developer Toolset 2.0. The Red Hat Developer Toolset 2.1 version of GCC does not introduce any new features or compatibility changes.

A.3. Changes in elfutils

Red Hat Developer Toolset 2.1 is distributed with **elfutils 0.157**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version and version 0.155 included in Red Hat Developer Toolset 2.0. Below is a comprehensive list of new features in this release.

A.3.1. Changes Since Red Hat Developer Toolset 2.0

The following features have been added since the release of elfutils included in Red Hat Developer Toolset 2.0:

- ▶ The **libdw** library now supports new functions **dwarf_getlocations()**, **dwarf_getlocation_attr()**, and **dwarf_getlocation_die()**.
- ▶ The **dwarf_getlocation()** function from the **libdw** library now supports new DWARF GNU extension opcodes **DW_OP_GNU_parameter_ref**, **DW_OP_GNU_convert**, **DW_OP_GNU_reinterpret**, **DW_OP_GNU_regval_type**, **DW_OP_GNU_entry_value**, **DW_OP_GNU_deref_type**, and **DW_OP_GNU_const_type**.
- ▶ The **dwfl_module_addrsym()** function from the **libdw** library now returns the **STB_GLOBAL** symbols in preference to **STB_WEAK**.
- ▶ The **dwfl_report_elf()** function from the **libdw** library now accepts an additional parameter. This new parameter can be used to configure whether the start address specified in another parameter is an absolute address or a bias relative to the **p_vaddr** pointer from the program header of the given ELF file.
- ▶ Functions **dwfl_module_getsymtab()**, **dwfl_module_getsym()**, **dwfl_module_addrsym()**, and **dwfl_module_addrname()** from the **libdw** library are now aware of the existence of MiniDebugInfo files embedded in main ELF files and consult them when looking for symbols.
- ▶ The **dwfl_core_file_report()** function from the **libdw** library now reads the chain from the **DT_DEBUG** entry before attempting to use other module discovery methods.
- ▶ The **eu-readelf** utility now understands core note type values **NT_ARM_TLS**, **NT_ARM_VFP**, **NT_ARM_HW_BREAK**, **NT_ARM_HW_WATCH**, **NT_SIGINFO**, and **NT_FILE** and converts them to a human-readable form. It also understands **.gdb_index** sections version 8.
- ▶ The **eu-readelf** utility now supports the following new command line options:
 - The **--elf-section** option can be used to inspect an embedded ELF file.
 - The **--unresolved-address-offsets** option or its short variant **-U** can be used to dump addresses in ELF files without resolving them to symbols.
 - The **--debug-dump=decodeline** option can be used to dump a **.debug_line** section in a human-readable, expanded format. In contrast, the original behavior of **--debug-dump=aranges** can now be achieved by using the **--debug-dump=decodedaranges** option, and the default formatting of the **.debug_aranges** section is now the same as in the ELF file.
- ▶ The **eu-addr2line** utility now supports a new command line option **--inlines** (and its short variant **-i**), which can be used to display all inlined subroutines of a particular address.

A.4. Changes in Git

Red Hat Developer Toolset 2.1 is distributed with **Git 1.8.4**, which provides a number of bug fixes and

feature enhancements over the Red Hat Enterprise Linux system version. Below is a comprehensive list of new features in this release.

A.4.1. Changes Since Red Hat Enterprise Linux 6.5

The features below have been added since the release of **Git** included in Red Hat Enterprise Linux 6.5.

General Changes in the Git Suite

- ▶ **Git** now detects the **DISCOVERY_ACROSS_FILESYSTEM** environment variable to decide whether to stop at a filesystem boundary when looking for the **.git** directory.
- ▶ The **GIT_ASKPASS** environment variable can now be used to tell commands from the **Git** suite that prompt for password to use an external program.
- ▶ The contents of the **/etc/mailname** file (if it exists) are now the default source of email hostnames for both the committer and the author.
- ▶ Conflict markers in the diff3 style now have a label.
- ▶ The **core.whitespace** configuration variable can now use a new rule: **tab-in-indent**. This rule treats a tab character at the beginning of a line as an error (useful for projects with space-only indents, for example, *Python* projects).
- ▶ The **core.eol** configuration variable is now the end-of-line authority over the working-tree files. It accepts the **lf**, **crlf**, or **native** values. The default is set to **native**.
- ▶ **Git** now reports any corrupt objects it finds.
- ▶ A dash ("-") is no longer a valid character at the beginning of a branch or tag name.
- ▶ The **pager.subcommand** configuration variable now allows for setting a different pager for each subcommand.
- ▶ The **core.abbrev** configuration variable can now be used to set the minimum length of abbreviated and unique object names.
- ▶ **Git** now recognizes **MATLAB** function and header patterns.
- ▶ **Git** now supports output coloring. A large number of configuration options have been added to allow for a separate tweaking of the look of various parts of outputs.
- ▶ **Git** commands now ignore the **.gitmodules** file if it is currently conflicted in a merge.
- ▶ The **-c name=value** command line option can now be used to override a configuration option.
- ▶ The **core.attributesfile** configuration variable can be set to a path of a file containing fallback default attributes. The default value is **/etc/gitattributes**.
- ▶ A new command, **git check-ignore**, has been added. The command helps to debug **.gitignore** files by showing the pattern on the basis of which individual files are included or excluded.
- ▶ The **Git** login shell for SSH accounts can now be extended by the user's own custom commands.
- ▶ A new configuration variable, **advice.statusHints**, is now available. It allows the user to disable (or enable) the hints about staging, unstaging, or adding files displayed by the **git status** command.
- ▶ A new credential helper has been added to access the GNOME Keyring and other services. Commands using some form of authentication can now use it to communicate with external programs. Refer to **gitcredentials(7)** for more information on configuring credential helpers.
- ▶ **Git** can now be told to interpret paths literally. The **--literal-paths** command line option serves that purpose. Setting the **GIT_LITERAL_PATHSPECS** environment variable to **1** has the same effect.
- ▶ The **.gitignore** file now accepts the **"**"** (double asterisk) recursive glob pattern.

- ▶ Regular expressions are now accepted in the **:/<string>** notation. Also, the expression is no longer bound to the start of the commit log message. Note that this change is not backward compatible.
- ▶ Custom messages are now issued to users who are attempting to interactively connect to the **git shell** when the interactive access is disabled.
- ▶ **Git** subcommands now honor the **i18n.*** configuration variables when handling text. See **git-config(1)** for more information on these settings.
- ▶ A considerable number of new scripts have been added to the **contrib/** directory.
- ▶ The **git fsck** and **git prune** commands are now reporting on their progress.
- ▶ The **http.useragent** configuration variable can now be set to lie to your firewall about the client used for **HTTP** connections.
- ▶ A new command, **git check-mailmap**, has been added. It can be used to map author and committer names and email addresses to their preferred names and addresses as specified in the **.mailmap** file.
- ▶ Users can now switch between commits with and without a submodule without re-cloning it due to improved plumbing — the **GIT_DIR** metainformation directory for submodules is set up inside a new submodule's directory.
- ▶ Whenever a date is expected as an argument, users can now specify a timezone designator format without seconds and with a colon between hours and minutes.
- ▶ The **GIT_PREFIX** environment variable is now available in the environments of processes spawned in configuration, so that they can check where in the tree the original command was invoked.
- ▶ A special path specification, **"/**, can be used to tell subcommands that work only with the current directory to work with the entire working tree.

The git am Command

- ▶ The **git am** command now accepts patches created by *Mercurial*.
- ▶ The **git am** command now sanitizes occurrences of "Re:" in the subject line.
- ▶ The following new command line options have been added to the **git am** command:
 - exclude=path**
Passed to the **git apply** command that is used in the background to apply the patch.
 - keep-cr**
Handles lines that contain a carriage return character (CR).

The git apply Command

- ▶ The **--verbose** (or **-v**) command line option now reports offsets in case a patch does not apply cleanly.

The git archive Command

- ▶ It is now possible to pipe the output of the **git archive** command to the **gzip** command. Do this by passing the **tar.gz** or **tgz** arguments to the **--format** command line option or by specifying either of these arguments in the **tar.format** configuration variable.

The git bisect Command

- ▶ The result of the **git bisect** command is now output as a comment in the bisect log file.

The `git blame` Command

- ▶ The **git blame** command now applies the configured *textconv* filters whenever possible. The *textconv* conversion allows **git blame** to present a human-readable output even when working with binary files.
- ▶ The following new command line options have been added to the **git blame** command:
 - abbrev[=*n*]**
Display *n*+1 hexadecimal digits instead of the default 7+1 as the abbreviated object name.
 - line-porcelain**
Use the easily parsable porcelain-formatted output, but show commit information for each line. This option implies the use of the **--porcelain** option. For a description of the porcelain format, see **git-blame(1)**.
 - show-email (or -e)**
Display email addresses instead of author names. This option is set to **off** by default.

The `git branch` Command

The following new command line options have been added to the **git branch** command:

- edit-description**
Open an editor to describe the purpose of the branch.
- list**
List existing branches. The option optionally accepts a globbing pattern.

The `git checkout` Command

The following new command line options have been added to the **git checkout** command:

- B *it***
A shorter version of running **git branch -f *it*** followed by **git checkout *it***. See **git-branch(1)** for more information.
- detach *commit***
A more user-friendly equivalent of the default behavior of **git *commit***.
- no-overwrite-ignore**
Do not overwrite untracked or ignored files during a merge.
- orphan *newbranch***
Creates a *newbranch* similarly to the **git checkout -b *newbranch*** command but with a detached root commit.

The `git check-attr` Command

- The **git check-attr** command now accepts relative paths passed from the command line.
- The following new command line options have been added to the **git check-attr** command:

--all (or -a)

List all the attributes of the specified path.

--cached

Look at the **.gitattributes** file in the index, and ignore the working tree.

The `git cherry-pick` Command

- The **git cherry-pick** command is now aware of commits that failed to apply.
- A range of commits can now be picked.
- A custom merge strategy can now be set. See below for a description of the new **--strategy** command line option.
- After resolving conflicts with **git cherry-pick -s**, sign-off lines are now implicitly left untouched in the log template.
- The following new command line options have been added to the **git cherry-pick** command:

--abort

Abort current operation and return to the state before the start of the sequence.

--allow-empty-message

Prevents the **git cherry-pick** command from failing when a commit is being picked that has no log message.

--continue

Continue the operation in progress. This option uses information found in the **.git/sequencer** file and can be used after conflict resolutions.

--strategy=strategy

Specifies the merge strategy for three-way merges. This option should only be used once. For information about the available strategies, see **git-merge(1)**.

The `git clean` Command

- The **git clean** command now provides an interactive mode, which can be invoked using the **--interactive** (or **-i**) command line option. For a description of subcommands available in the interactive mode, see **git-clean(1)**.
- A new command line option, **--exclude=pattern** (or **-e pattern**), has been added. This option allows the user to define rules that specify which files to ignore (in addition to what is defined in the **.gitignore** and **GIT_DIR/info/exclude** files).

The `git clone` Command

- A new command line option, **--config key=value** (or **-c key=value**), is now available. This option sets the specified configuration variable to be in effect at the time of the initial checkout. The syntax expected is the same as what is used with the **git config** command.

The git commit Command

- ▶ The **git commit** command now accepts arguments in the form of relative paths that point outside of the current directory.
- ▶ Aborting interactive mode (initialized by the **--interactive** command line option) now discards all changes that were made in it.
- ▶ It is now possible to set the **commit.cleanup** configuration variable to **whitespace**, which is equivalent to running the **git commit** command with the **--cleanup=whitespace** command line option. For a description of all available cleanup modes, see **git-commit(1)**.
- ▶ The following new command line options have been added to the **git commit** command:
 - fixup=commit**
Prefix the commit line with “fixup! ”. This option is used when running the **git rebase --autosquash** command.
 - gpg-sign[=keyid] (or -S[keyid])**
Sign the commit using the specified GPG key.
 - no-edit**
Do not launch an editor for the commit message. When used with the **--amend** command line option, it only updates the commit tree, not the commit message.
 - patch (or -p)**
Select patch hunks for committing using the interactive interface. Refer to **git-add(1)** for more information about the interactive mode.
 - squash=commit**
Prefix the commit line with “squash! ”. This option is used when running the **git rebase -autosquash** command.

The git count-objects Command

- ▶ The verbose output, invoked by passing the **--verbose** (or **-v**) command line option to the **git count-objects** command, now reports on various garbage in the **.git** directory. Refer to **git-count-objects(1)** for a more detailed description of the verbose output.
- ▶ A new command line option, **--human-readable** (or **-H**), has been added. This option makes **git count-objects** show sizes in a human-readable format.

The git cvsimport Command

- ▶ A new command line option, **-R**, has been added to **git cvsimport**. This option creates a mapping between **CVS** revisions and **Git** commits.
- ▶ Users of the **git cvsimport** command can now use the following configuration variables to enable certain command line options by default:

| Configuration Variable | Command Line Option | Description |
|-------------------------------|------------------------------|---|
| cvsimport.authorsFiles | -A file_name | Maps CVS user names to full names and email addresses. |
| cvsimport.ignorePaths | -S regular_expression | Skips paths that match a certain regular expression. |
| cvsimport.mergeRegex | -M regular_expression | Parses commit messages in order to |

| | |
|------------------------------------|--|
| cvsimport.trackRevisions -R | detect merge commits. |
| | Creates a mapping between CVS revisions and Git commits. |

The git cvsserver Command

- The **git cvsserver** can now use **pserver** to offer **CVS** access. Refer to **git-cvsserver(1)** for information on how to configure this.

The git daemon Command

- The **git daemon** server can now listen on multiple addresses.
- The daemon now allows for external hooks via the **--access-hook=path** command line option. This option executes an external command specified by **path** every time a client connects to the daemon.

The git describe Command

- A new command line option, **--first-parent**, has been added. This option instructs the **git describe** command to only follow the first parent commit in case of merge commits.

The git diff Command

- The **--submodule** command line option is now able to notice 'dirty' submodules. These are submodules that contain modified or untracked files.
- The regular expression for function headers now detects visibility modifiers (public, private, or protected) in the *PHP5* programming languages.
- Functions and subroutines of the *Fortran*, *Pascal*, and *Perl* programming languages are now recognized.
- A new configuration variable, **diff.submodule**, is now available for setting the default format used for showing differences in submodules using the **--submodule** command line option. The option can have two values: **log** or **short**. For a detailed description of the two values, see **git-diff(1)**. The default value is **short**.
- A new configuration variable, **diff.noprefix**, is now available for setting an implicit behavior identical to what the **--no-prefix** command line option does. If set, no source or destination prefix is shown.
- A new configuration variable, **diff.context**, is now available for setting the number of lines shown above and below the modified line or lines. The default value is **3**. The **--unified** (or **-U**) option overrides this setting.
- A new configuration variable, **diff.algorithm**, is now available. This variable can be used to specify the algorithm used to display the **git diff** output. It accepts the following arguments: **default** or **myers**, **minimal**, **patience**, and **histogram**. For a description of the different algorithms, refer to **git-diff(1)**. The default value is **default** or **myers**.
- A new configuration variable, **diff.ignoresubmodules**, is now available. This variable can be used to ignore submodules when generating diffs.
- The following new command line options have been added to the **git diff** command:
 - function-context (or -w)**
Displays entire functions in the context of changes.
 - histogram**

Generates output using the *histogram* algorithm. This algorithm supports low-occurrence common elements and offers better performance than the default.

--ignore-blank-lines

Tells **git diff** to ignore changes that consist only of additions or removals of blank lines.

--minimal

Generates output using the *minimal* algorithm. This algorithm attempts to minimize the output at the cost of extra time spent.

--stat-count

Limits the length of output to the number of lines specified by the option. Lines exceeding the limit are substituted with "...".

--word-diff=mode

Shows added or removed words. This option accepts **color**, **plain**, **porcelain**, and **none** as modes. For a description of the available modes, refer to **git-diff(1)**. The default is **plain**.

The git fetch Command

- The **git fetch** command now honors the **transfer.fsckobjects** configuration variable. If set, the command validates received objects.
- A new command line option, **--recurse-submodules[=argument]**, has been added. This option controls whether to fetch new commits from submodules too. This option accepts the **yes**, **no**, or **on-demand** arguments. If passed without an argument, it assumes **yes**. For a description of the arguments, see **git-fetch(1)**.

The git format-patch Command

- The **git format-patch** command now accepts the **--subject-prefix** command line option with no argument. This option substitutes the standard "[PATCH]" prefix in the subject line with a new prefix (or none) specified on the command line.
- The **format.coverLetter** configuration variable may now be used to turn the **--cover-letter** command line option **on** or **off**. The **format.coverLetter** may further be set to **auto** to only be used with sets of two or more commits.
- The following new command line options have been added to the **git format-patch** command:
 - from[=whom]**
Set the "From:" field to the specified person. Not to be used when preparing patches for the **git send-mail** command. See **git-format-patch(1)** for more information about the behavior of this option.
 - notes[=ref]**
Add notes to the formatted patch. Refer to **git-notes(1)** for more information on the way **Git** handles these notes.
 - [no-]signature**
Alter the email signature used in the output messages. The **format.signature** configuration variable can be used to set the default signature. When neither the command line option nor the configuration variable is specified, the **Git** version number is used as the

signature. The **--no-signature** variant of the option can be used to prevent the inclusion of any signature.

--quiet (or -q)

Do not show the names of the generated patch files.

--reroll-count=*n* (or -v *n*)

Output a set of patches with the *n* version string prepended to their names and the version mentioned in the subject.

The git for-each-ref Command

- The **git for-each-ref** command now recognizes the following new fields for extracting values from referenced objects into the resulting output: **%(symref)**, **%(symref:short)**, **%(flag)**, **%(objectname:short)**, **%(contents:subject)**, **%(contents:body)**, and **%(contents:signature)**.

The git grep Command

- Functions and subroutines of the *Fortran*, *Pascal*, and *Perl* programming languages are now recognized.
- It is now possible to pass the - (dash) argument to the **-f** command line option to instruct it to read patterns from the standard input.
- The following new command line options have been added to the **git grep** command:

--break

Delimit matches from different files with an empty line.

--color=[*when*]

Color the matches returned by **git grep**. This option accepts the **auto**, **always**, or **never** arguments. The default setting is **always**.

--exclude-standard

Do not search in files listed in the **.gitignore** file when using the **--no-index** option.

--function-context (or -W)

Show the entire function in which the match was found.

--heading

Print the filename above the matches from that file instead of showing the file information for each match separately.

--no-index

In the current directory, search also for matches in files that are not managed by **Git**.

--open-files-in-pager [*pager*] (or -O [*pager*])

Instructs **git grep** to open files with matching regular expressions in the user's pager or editor of choice.

-P

Instructs **git grep** to expect and accept *Perl*-compatible regular expressions.

--untracked

In the current tree, search in both tracked and untracked files.

The git hash-object Command

- The **--stdin-paths** command line option can now be used together with the **--no-filter** option. This results in the **git hash-object** command reading filenames from the standard input and hashing the contents as is — disregarding any filters that may have been applied by the attributes mechanism.

The git help Command

- The **--web** (or **-w**) command line option now works with the **Chrome** and **Chromium** web browsers. The default web browser can be specified using the **help.browser** or **web.browser** configuration variables. If none of these variables is defined, the **git web--browse** script attempts to choose a browser automatically.
- A new command line option, **--guides** (or **-g**), has been added. This option shows a list of useful guides.

The git imap-send Command

- The **git imap-send** command now supports *CRAM-MD5* as a valid method for authentication with an *IMAP* server.

The git init Command

- A new configuration variable, **init.templateDir**, is now recognized. Similarly to the **--template** command line option and the **GIT_TEMPLATE_DIR** environment variable, it specifies the directory containing files and directories that will be copied to the **.git** directory (or **GIT_DIR**).
- A new command line option, **--separate-git-dir=gitdirectory**, has been added. This option forces **git init** to initialize the repository in the specified location. A symbolic link is created in place of the **.git** directory, pointing to the separately placed repository.

The git log Command

- The **--follow** command line option now follows copies as well as renames when listing the history of a file.
- The **git log** commands are now able to interpret paths specified using globs.
- The **--format=format** option is now able to interpret the **%g[nNeE]** specifiers when showing information from *reflogs*. For a description of the specifiers, see **git-log(1)**.
- The **--decorate** command line option now shows commit information in colors. It is now also capable of highlighting grafted and replaced commits.
- The following new command line options have been added to the **git log** command:
 - author-date-order**
Sort the commits listed based on the author timestamp. No parents are shown before all of their children are listed.
 - cherry**
Display two diverged branches, and omit any changes they share. This is equivalent to the following combination of options: **--right-only --cherry-mark --no-merges**.

--cherry-mark

Display two diverged branches, and mark any changes they share with the “=” (equal sign) character.

--children

List also the children of commits.

-Gpattern

Only match changes that contain *pattern*.

--grep-reflog=pattern

Limit the output to commits whose reflog messages match *pattern*. If multiple **--grep-reflog** options are used, commits with reflog messages matching any of the patterns specified are listed. This option should only be used together with the **--walk-reflogs** option.

-Lstart,end:file

Show the changes that the line specified by the range within the *file* went through. Either end of the range may be specified by absolute line numbers or regular expressions. The end of the range may also be an offset that specifies a number of lines before or after *start*.

--notes

Replaces the **--show-notes** option. This option shows the notes annotating the commits (commit log messages).

--show-signature

Pass the signature of a commit object to the **gpg --verify** command to check its validity, and show the output.

--use-mailmap

Rewrite the names and email addresses of committers in the output using the **mailmap** file.

The git ls-remote Command

- ▶ A new command line option, **--exit-code**, has been added. This option returns 2 as the command's exit code when no matching remote refs are found.
- ▶ A new command line option, **--help**, has been added. This option displays basic information about the syntax and options of the **git ls-remote** command.

The git merge Command

- ▶ The **--log[n]** command line option may now be optionally given a numeric parameter. The parameter specifies the number of commits merged from which one-line descriptions will be used in the log message in addition to branch names.
- ▶ It is now possible to use the **git merge -** command, where **-** (dash) denotes the name of the previous branch.
- ▶ The **merge.ff** configuration variable can now be set to **yes**, **no**, or **only** to specify the default behaviour, which corresponds to the **--ff**, **--no-ff**, and **--ff-only** command line options respectively. For a description of these options, see **git-merge(1)**.

- The following new command line options have been added to the **git merge** command:

--abort

Abort the current conflict-resolution process, and attempt to go back to the pre-merge state. It may not be possible to go back if there were uncommitted changes when the merge was started.

--edit (or -e)

Open an editor to edit the automatically generated merge message before committing a successful merge.

--no-overwrite-ignore

Do not overwrite untracked or ignored files during a merge.

--[no-]verify-signature

Verify the GPG signatures of the commits that are being merged. Abort the merge if the verification fails. The **--no-verify-signature** variant of the option instructs the **git merge** command not to verify commit signatures.

The git mergetool Command

- The **git mergetool** command now instructs the **vim** or **gvim** editors to show the three-way diff by default. The old behavior can be achieved by specifying **vimdiff2** or **gvimdiff2** as the tool to be used.

The git merge-base Command

- A new command line option, **--is-ancestor commit commit**, has been added. This option returns the 0 exit status code if the first **commit** is an ancestor of the second **commit** and 1 if not.

The git merge-file Command

- A custom conflict-marker size can now be used. The custom size can be set using the **--marker-size** command line option or the **conflict-marker-size** configuration attribute.
- The union merge behaviour can now be used. This means that **git merge-file** resolves conflicts by concatenating lines from both sides of the merge.

The git notes Command

- A new subcommand, **git notes merge**, is now available. This subcommand attempts to merge notes accompanying the specified object into the current notes ref.

The git notes prune Subcommand

The following new command line options have been added to the **git notes prune** subcommand:

--dry-run (or -n)

Does not remove anything. Instead, it only lists the names of objects whose notes would be pruned.

--verbose (or -v)

Makes the **prune** subcommand provide more information when merging notes and report all

names of objects whose notes are removed during the pruning.

The git pull Command

- The **git pull** command can now pull an annotated or signed tag instead of a branch tip. The GPG signature of the tag is then made a part of the merge commit.
- The **pull.rebase** configuration variable is now recognized. It makes **git pull** act as if the **--rebase** command line option was passed. The variable serves as a fallback setting in case the per-branch **branch.name.rebase** variable is not set. Note that this may be a potentially dangerous setting. Refer to **git-rebase(1)** for more information about rebasing and to **git-config(1)** for a description of the **pull.rebase** variable.
- The **--verbose** (or **-v**) and **--quiet** (or **-q**) command line options are now passed along to the **git rebase** command when **git pull --rebase** is run or when the corresponding **branch.name.rebase** configuration variable is set.
- A new command line option, **--[no-]verify-signature**, has been added. This option verifies the GPG signatures of the commits that are being merged and aborts the merge if the verification fails. The **--no-verify-signature** variant of the option instructs the **git pull** command not to verify the GPG signatures of the commits.

The git push Command

- A new “pre-push” hook is available. If the hook exits with a non-zero status code, it will prevent the **git push** command from performing any operations. To enable the hook, pass the **--verify** command line option to **git push**. Use the **--no-verify** option to bypass the hook. For a more detailed description of the hook, refer to **githooks(5)**.
- The **--force** (or **-f**) command line option is now required for updating remote tags. Note that forcing a push can result in a loss of commits in the remote repository if not used with care. For a more detailed description of this option, refer to **git-push(1)**.
- The **push.default** configuration variable has a new valid value of **upstream**, which is synonymous to the now deprecated **tracking** value. For a description of the available **push.default** values, refer to **git-config(1)**.
- The following new command line options have been added to the **git push** command:
 - follow-tags**
Push annotated tags plus all the refs that would be pushed normally.
 - recurse-submodules=check|on-demand**
Do not push if any submodule is ahead of its remote HEAD. See **git-push(1)** for an explanation of the two possible arguments.

The git read-tree Command

- A new command line option, **--dry-run** (or **-n**), has been added. This option checks if a merge would complete without errors when the **git read-tree** command is passed the **-m** option.

The git rebase Command

- The **rebase.autosquash** configuration option is now honored by the **git rebase --interactive** command. The **--interactive** (or **-o**) command line option prepares a list of the

commits to be rebased and lets the user modify that list before the rebasing takes place.

- The **git rebase** command now updates reflog for HEAD even when it finishes normally or is ended with the **--abort** command line option.
- It is now possible to pass SHA-1 hashes as object names when using the **--autosquash** command line option.
- The **--strategy-option=option** (or **-X option**) command line option can now be used along with the **--strategy=strategy** (or **-s strategy**) option to pass options to the selected merge strategy. Refer to **git-rebase(1)** for a description of the available strategies.
- The following new command line options have been added to the **git rebase** command:

--edit-todo

Edit the instructions in an editor when using the interactive mode (invoked by the **--interactive** (or **-i**) option).

--[no-]autostash

Temporarily stash uncommitted changes before the rebasing begins instead of refusing to run. The **--no-autostash** variant of the option may be used to prevent autostashing when the **rebase.autostash** configuration variable is set.

The git remote Command

- A new subcommand, **git remote set-branches**, is available. This subcommand modifies the list of branches tracked by the specified remote.

The git replace Command

- The **--delete** (or **-d**) command line option now interprets its **object** argument as the extended SHA-1 syntax. The **object** may now be specified, for example, as **HEAD~NUM**.

The git rerere Command

- A new subcommand, **git rerere remaining**, is now available. It displays conflicts that have not been automatically resolved. This includes paths that the **git rerere** command cannot track, such as submodules.
- The **git rerere forget** subcommand now requires a path to be specified where any conflict resolutions are to be reset.

The git reset Command

- A new command line option, **--keep**, has been added. This option discards index entries and commits in the working tree but preserves local changes.

The git revert Command

- A new command line option, **--strategy=strategy**, has been added. This option specifies the merge strategy for three-way merges. This option should only be used once. For information about the available strategies, see **git-merge(1)**.

The git send-email Command

- The **git send-email** command is now able to parse and use the “To:” field from its input files.

- ▶ The value of the **sendmail.aliasfile** configuration variable may now begin with the “~” (tilde) character to denote the user's home directory.
- ▶ The **sendmail.annotate** configuration variable may now be used to set the default behavior that is otherwise controlled by the **--annotate** command line option. If set to **true**, each patch that is about to be sent is offered for review and editing.
- ▶ Patches sent with the **--in-reply-to** command line option are now replies to the new cover letter that begins the new thread.
- ▶ The following new command line options have been added to the **git send-email** command:
 - help (or -h)**
Display basic information about the syntax and options of the **git send-email** command.
 - to-cmd=command**
This option is similar to the **--cc-cmd=command** option and can be used to read the list of recipients from the output of the **command**.

The git show Command

- ▶ The **git show --do-walk -n** command now has a more compact equivalent, **git show -n**. Refer to **git-rev-list(1)** for more information on object traversal.
- ▶ The **--no-patch** command line option is a new synonym for the **-s** option. This option suppresses diff output. It can also be used to cancel the effect of the **--patch** option.

The git show-ref Command

- ▶ The **--head** command line option now always means the “HEAD” reference, regardless of filtering.

The git stash Command

- ▶ A new command line option, **--include-untracked** (or **-u**), has been added. This option stashes untracked files as well, which allows for them to be cleaned later by the **git clean** command.

The git status Command

- ▶ The **git status** command now detects submodules with uncommitted local changes in their working tree.
- ▶ Users can now change the value of the **status.short** configuration variable to **true** to configure the **git status** command to use the **--short** command line option by default. Similarly, it is now possible to change the value of the **status.branch** configuration variable to **true** to run **git status** with the **--branch** option by default.
- ▶ The following new command line options have been added to the **git status** command:
 - branch (or -b)**
Displays the name of the current branch even when the **--short** (or **-s**) option is used.
 - ignore-submodules[=condition]**
This option accepts **none**, **untracked**, **dirty**, or **all** as an optional argument and can be used to ignore certain changes to submodules.
 - ignored**
Displays ignored files.

The git submodule Command

- ▶ Users are now allowed to create shallow clones of submodule repositories.
- ▶ Users can now run **git submodule** commands from any part of the working tree.
- ▶ The **git submodule sync** command now also updates information for submodules that have not been checked out.
- ▶ The **git submodule update** command no longer terminates immediately when it encounters an error. Instead, it updates the remaining registered submodules and displays the error message on exit.
- ▶ The **git submodule update** command can now be configured to execute an arbitrary shell command. To enable this behavior for a particular submodule, change the value of the **submodule.submodule_name.update** configuration variable to any shell command that accepts exactly one command line argument.
- ▶ Users can now set the value of the **submodule.submodule_name.update** configuration variable to **none** to prevent a particular submodule from being updated by default.
- ▶ A new command, **git submodule deinit**, has been added. This command unregisters selected submodules and serves as a counterpart to the **git submodule init** command.
- ▶ A new command line option, **--force** (or **-f**), has been added to **git submodule**. This option can be used to discard local changes.
- ▶ A new command line option, **--name**, has been added to **git submodule**. This option can be used to specify the name of a new submodule.
- ▶ It is now possible to apply the **--recursive** command line option to the **sync** command of **git submodule**. This option makes **sync** traverse nested submodules recursively.

The git tag Command

- ▶ A new command line option, **--cleanup=mode**, has been added. This option controls the way whitespace is handled in tag messages. It accepts the **verbatim**, **whitespace**, or **strip** arguments. See **git-tag(1)** for a description of what the available modes do.

The gitweb Interface

- ▶ The configuration of the gitweb interface can now be changed after each request.
- ▶ The **/etc/gitweb-common.conf** configuration file is now read if it exists.
- ▶ The gitweb interface is now able to show diffs.

Other new subcommands and exotic option combinations

The following subcommands and exotic option combinations have been added:

git remote -t branch --mirror=fetch

This is now a legal combination of options as the mirroring in the **fetch** mode no longer mirrors everything.

git rev-list A..B --ancestry-path

The newly available **--ancestry-path** option excludes commits that are not descendants of commit **A**.

git rev-list --count --cherry-mark

This combination now produces a more useful output as it counts the cherry-picked commits separately.

A.5. Changes in GDB

Red Hat Developer Toolset 2.1 is distributed with **GDB 7.6.1**, which provides a number of bug fixes over the Red Hat Enterprise Linux system version and version 7.6 included in Red Hat Developer Toolset 2.0. Below is a comprehensive list of new features in this release.

A.5.1. Changes Since Red Hat Developer Toolset 2.0

The following features have been added since the release of GDB included in Red Hat Developer Toolset 2.0:

- It is now possible to use the ***ClassName::MemberVariable*** expression to access non-static member variables of a particular class in the current scope of implicit parameter **this**.
- GDB now automatically loads files with the **.gdb** and **.py** file extension located in the **/opt/rh/devtoolset-2/root/etc/gdbinit.d/** directory on startup.

A.6. Changes in Valgrind

Red Hat Developer Toolset 2.1 is distributed with **Valgrind 3.8.1**, which provides a number of bug fixes over the version included in Red Hat Developer Toolset 2.0. Below is a comprehensive list of new features in this release.

A.6.1. Changes Since Red Hat Developer Toolset 2.0

The following features have been added since the release of Valgrind included in Red Hat Developer Toolset 2.0:

- Manual pages **cg_diff(1)**, **cg_merge(1)**, **valgrind-listener(1)**, and **vgdb(1)** have been added.
- Support for **PTRACE_GET_THREAD_AREA** and **PTRACE_SET_THREAD_AREA** has been added.
- Valgrind now supports **AF_BLUETOOTH** sockets.
- Valgrind now supports the **SSE4 MOVNTDQA** instruction.

A.7. Changes in OProfile

Red Hat Developer Toolset 2.1 is distributed with **OProfile 0.9.8**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version. Below is a comprehensive list of new features in this release.

A.7.1. Changes Since Red Hat Developer Toolset 2.0

The following features have been added since the release of OProfile included in Red Hat Developer Toolset 2.0:

- OProfile now supports the Intel Core i3, Intel Core i5, Intel Core i7, Intel Xeon Processor E3, Intel Pentium, and Intel Celeron processors formerly code-named Haswell.
- Users can now use named unit masks for Intel processors.

Changes in Version 2.0

The sections below document features and compatibility changes introduced in Red Hat Developer Toolset 2.0.

B.1. Changes in Eclipse

Red Hat Developer Toolset 2.0 is distributed with **Eclipse 4.3** and a wide selection of plug-ins from the Eclipse Foundation's 2013 Eclipse 4.3 “Kepler” release, and provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version.

Below is a comprehensive list of new features and compatibility changes in this release. For details on how to use these new features, see the built-in Eclipse documentation.

B.1.1. Changes Since Red Hat Enterprise Linux 6.4

The following features have been added since the release of the Eclipse Platform and native plug-ins included in Red Hat Enterprise Linux 6.4:

- ▶ The **Eclipse Platform** has been updated from version 3.6 to 4.3. This new major release features a more flexible user interface based on Eclipse Modeling Framework (EMF), CSS-based themes, dependency injection, and more.
- ▶ The **Eclipse C/C++ Development Toolkit** (CDT) has been updated from version 7.0 to 8.2. This new major release includes:
 - a greatly enhanced indexer, both in terms of capabilities and speed;
 - improved GNU Debugger integration;
 - support for GNU Debugger's pretty-printing;
 - multi-process debugging;
 - project-less debugging;
 - enhanced refactoring;
 - code-checking capabilities;
 - an improved and more polished user interface.
- ▶ **Mylyn**, a set of task and application life cycle management plug-ins, has been updated from version 3.4 to 3.9. This update includes:
 - an enhanced task-focused interface and task editing;
 - a new Jenkins/Hudson connector;
 - enhanced Bugzilla and Trac connectors;
 - EPUB authoring tools.
- ▶ A new **EGit** plug-in has been added. This plug-in includes an Eclipse Team provider based on JGit, a Git implementation written entirely in Java, and adds support for the Git revision control system to the Eclipse user interface by introducing the **History** and **Synchronize** views, **Compare** and **Quick Diff** menu items, and various wizards to make it easier for developers to use this plug-in.
- ▶ New **GCov** integration has been added to allow the user to visualize GCov output in both summarized form and in file editors. As well, graphing capabilities for the data have been added.
- ▶ New **GProf** integration has been added to allow the user to view profiling information such as execution time and call graphs. As well, graphing capabilities for the data have been added.
- ▶ A new **SystemTap** integration plug-in has been added. This plug-in includes:
 - an editor for the **.stp** files with the autocomplete feature;
 - the **Probe** view with a list of probes that are available on the system;

- the **Function** view with a list of functions that are available on the system and can be used in **.stp** scripts.

As well, this plug-in includes integration for running SystemTap scripts and viewing the results in a textual, tabular, and graphical manner. Note that the result sets are updated in a near-runtime way, which allows the user to use this plug-in for longer-running monitoring tasks.

- New kernel **perf** tool integration has been added. This plug-in uses the performance counters subsystem of the Linux kernel to profile applications, makes it easier to analyze the results by hyperlinking to the sources in the workspace projects, simplifies the **perf** tool configuration by selecting the counters to be used, and allows the user to run this tool remotely.
- A unified profiling launcher has been added to provide a single method to launch profiling. It allows the user to select a profiling category (**Memory**, **Timing**, or **Coverage**) and back ends for this category (such as **OProfile**, **perf**, or **GProf** for **Timing**).
- The C/C++ documentation plug-in has been enhanced to recognize and use **gtk-doc** generated documentation and to display it in the Eclipse Help Center.
- The **OProfile** plug-in has been enhanced to support root privilege authentication through **polkit**. This feature is configured automatically.
- The **Valgrind** plug-in now supports the **Helgrind** tool, which can be used to detect synchronization problems such as race conditions, deadlocks caused by incorrect locking order, or misuse of the POSIX **pthread**s API. When a problem is detected, the plug-in displays error markers on the corresponding lines in the source code.

B.2. Changes in GCC

Red Hat Developer Toolset 2.0 is distributed with **GCC 4.8**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version and the version included in Red Hat Developer Toolset 1.1. Below is a comprehensive list of new features and compatibility changes in this release.

B.2.1. Changes Since Red Hat Developer Toolset 1.1

The following features have been added since the release of GCC included in Red Hat Developer Toolset 1.1.

B.2.1.1. Caveats

Aggressive Loop Optimizations

The loop optimizer of GCC has been improved to use language constraints in order to derive bounds for the number of iterations of a loop. The bounds are then used as a guide to loop unrolling, peeling, and loop exit test optimizations.

The optimizations assume that the loop code does not invoke undefined behavior by, for example, causing signed integer overflows or making out-of-bound array accesses. For example, consider the following code fragment:

```

unsigned int foo()
{
    unsigned int data_data[128];

    for (int fd = 0; fd < 128; ++fd)
        data_data[fd] = fd * (0x02000001); // error

    return data_data[0];
}

```

When the value of the `fd` variable is 64 or above, the `fd * 0x02000001` operation overflows, which is invalid in both C and C++ for signed integers. In the example above, GCC may generate incorrect code or enter an infinite loop.

To fix this error, use the appropriate casts when converting between signed and unsigned types to avoid overflows, for instance:

```

data_data[fd] = (uint32_t) fd * (0x02000001U); // ok

```

If necessary, this optimization can be turned off by using the new command line option **-fno-aggressive-loop-optimizations**.

B.2.1.2. General Improvements and Changes

New Local Register Allocator

GCC 4.8 features a new *Local Register Allocator* (LRA), which replaces the 26-year old reload pass and improves the quality of generated code. The new local register allocator is meant to be simpler, easier to debug, and does a better job of register allocation.

AddressSanitizer

A fast memory error detector called *AddressSanitizer* has been added and can be enabled by using the **-fsanitize=address** command line option. It augments memory access instructions in order to detect use-after-free and out-of-bound accesses to objects on the heap.

ThreadSanitizer

A fast data race detector called *ThreadSanitizer* has been added in GCC 4.8. The option to enable this feature is **-fsanitize=thread**.

Compiling Extremely Large Functions

Many scalability bottlenecks have been removed from GCC optimization passes. As a consequence, it is now possible to compile extremely large functions with smaller memory consumption in less time.

New -Og Optimization Level

A new general optimization level, **-Og**, has been introduced. This optimization level addresses the need for fast compilation and a superior debugging experience while providing a reasonable level of runtime performance. Overall, the development experience should be better than the default optimization level **-O0**.

Caret Diagnostic Messages

The diagnostic messages of GCC, which display a line of source code, now also show a caret that indicates the column where the problem was detected. For example:

```
fred.cc:4:15: fatal error: foo: No such file or directory
#include <foo>
          ^
compilation terminated.
```

New -fira-hoist-pressure Option

A new command line option, **-fira-hoist-pressure**, has been added. This option uses the register allocator to help decide when it is worthwhile to move expressions out of loops. It can reduce the size of the compiler code, but it slows down the compiler. This option is enabled by default at **-Os**.

New -fopt-info Option

A new command line option, **-fopt-info**, has been added. This option controls printing information about the effects of particular optimization passes, and takes the following form:

```
-fopt-info [-info] [=file_name]
```

The **info** part of the option controls what is printed. Replace it with **optimized** to print information when optimization takes place, **missed** to print information when optimization does not take place, **note** to print more verbose information, or **optall** to print everything.

Replace **file_name** with the name of the file in which you want the information to be written. If you omit this part of the option, GCC writes the information to the standard error output stream.

For example, to display a list of optimizations that were enabled by the **-O2** option but had no effect when compiling a file named **foo.c**, type:

```
gcc -O2 -fopt-info-missed foo.c
```

New -floopt-nest-optimize Option

A new command line option, **-floopt-nest-optimize**, has been added. This option enables an experimental ISL-based loop nest optimizer, a generic loop nest optimizer that is based on the Pluto optimization algorithms and that calculates a loop structure optimized for data-locality and parallelism. For more information about this optimizer, see <http://pluto-compiler.sourceforge.net>.

Hot and Cold Attributes on Labels

The hot and cold function attributes can now also be applied to labels. Hot labels tell the compiler that the execution path following the label is more likely than any other execution path, and cold labels convey the opposite meaning. These attributes can be used in cases where **__builtin_expect** cannot be used, for instance with a computed **goto** or **asm goto**.

B.2.1.3. Debugging Enhancements

DWARF4

DWARF4 is now used as the default debugging data format when generating debugging information. To get the maximum benefit from this new debugging representation, use the latest version of **Valgrind**, **elfutils**, and **GDB** included in this release.

New -gsplit-dwarf Option

A new command line option, **-gsplit-dwarf**, has been added. This option tells the compiler driver to separate as much DWARF debugging information as possible into a separate output file with the **.dwo**

file extension, and allows the build system to avoid linking files with debugging information.

In order to be useful, this option requires a debugger capable of reading **.dwo** files, such as the version of **GDB** included in Red Hat Developer Toolset 2.0.



Note

elfutils, **SystemTap**, and **Valgrind** do *not* support the **.dwo** files.

B.2.1.4. C++ Changes

Experimental C++ Features from an Upcoming Standard

g++ now supports a new command line option, **-std=c++1y**. This option can be used for experimentation with features proposed for the next revision of the standard that is expected around 2014. Currently, the only difference from **-std=c++11** is support for return type deduction in normal functions as proposed in [N3386](#).

New **thread_local** Keyword

g++ now implements the C++11 **thread_local** keyword. In comparison with the GNU **__thread** keyword, **thread_local** allows dynamic initialization and destruction semantics.

The use of the **thread_local** keyword has currently one important limitation: when the **dlclose()** function is used to unload a dynamically loaded DSO that contains the definition of a **thread_local** object, the **thread_local** object is destroyed, its destructor is called and the DSO is unmapped from the address space of the process. If a thread in the process tries to access the **thread_local** object after this, the program may terminate unexpectedly. As a result, the programmer may have to take extra care to ensure that **thread_local** objects in a DSO are not referred after it has been unloaded.

See also the next item for dynamic initialization issues.

Dynamic Initialization of Thread-local Variables

The C++11 and OpenMP standards allow thread-local and thread-private variables to have dynamic (that is, runtime) initialization. To support this, any use of such a variable goes through a wrapper function that performs necessary initialization.

When the use and definition of the variable are in the same translation unit, this overhead can be optimized away, but when the use is in a different translation unit, there is significant overhead even if the variable does not actually need dynamic initialization. If the programmer can be sure that no use of the variable in a non-defining translation unit needs to trigger dynamic initialization (either because the variable is statically initialized, or a use of the variable in the defining translation unit will be executed before any uses in another translation unit), they can avoid this overhead by using the new **-fno-extern-tls-init** option.

By default, **g++** uses the **-fextern-tls-init** option.

C++11 Attribute Syntax

g++ now implements the C++11 attribute syntax, for example:

```
[[noreturn]] void f();
```

C++11 Alignment Specifier

g++ now implements the C++11 alignment specifier, for example:

```
alignas(double) int i;
```

B.2.1.5. Fortran Changes

B.2.1.5.1. Caveats

The version of module files (the **.mod** files) has been incremented. Fortran modules compiled by earlier GCC versions have to be recompiled when they are used by files compiled with GCC 4.8, as this version of GCC is not able to read **.mod** files created by earlier versions; attempting to do so fails with an error message.



Note

The ABI of the produced assembler data itself has not changed; object files and libraries are fully compatible with older versions except as noted in [Section B.2.1.5.2, “ABI Compatibility”](#).

B.2.1.5.2. ABI Compatibility

Some internal names used in the assembler or object file have changed for symbols declared in the specification part of a module. If an affected module — or a file using it via use association — is recompiled, the module and all files which directly use such symbols have to be recompiled as well. This change only affects the following kind of module symbols:

- ▶ *Procedure pointers*. Note that C-interoperable function pointers (**type(c_funptr)**) are not affected, nor are procedure-pointer components.
- ▶ *Deferred-length character strings*.

B.2.1.5.3. Other Changes

BACKTRACE Intrinsic

A new intrinsic subroutine, **BACKTRACE**, has been added. This subroutine shows a backtrace at an arbitrary place in user code, program execution continues normally afterwards.

Floating Point Numbers with “q” as Exponential

Reading floating point numbers that use **q** for the exponential (such as **4.0q0**) is now supported as a vendor extension for better compatibility with old data files. It is strongly recommended to use the equivalent but standard conforming **e** (such as **4.0e0**) for I/O.

For Fortran source code, consider replacing the **q** in floating-point literals by a kind parameter (such as **4.0e0_qp** with a suitable **qp**). Note that — in Fortran source code — replacing **q** with a simple **e** is not equivalent.

GFORTRAN_TMPDIR Environment Variable

The **GFORTRAN_TMPDIR** environment variable for specifying a non-default directory for files opened with **STATUS="SCRATCH"**, is not used anymore. Instead, **gfortran** checks the POSIX/GNU standard **TMPDIR** environment variable and if **TMPDIR** is not defined, **gfortran** falls back to other methods to determine the directory for temporary files as documented in the user manual.

Fortran 2003

Support for unlimited polymorphic variables (**CLASS(*)**) has been added. Non-constant character lengths are not yet supported.

TS 29113

Assumed types (**TYPE(*)**) are now supported.

Experimental support for assumed-rank arrays (**dimension(..)**) has been added. Note that at the moment, the **gfortran** array descriptor is used, which is different from the array descriptor defined in *TS 29113*. For more information, see the header file of **gfortran** or use the **Chasm** language interoperability tools.

B.2.1.6. x86-specific Improvements

New Instructions

GCC 4.8 has added support for the Intel **FXSR**, **XSAVE**, and **XSAVEOPT** instructions. Corresponding intrinsics and built-in functions can now be enabled by using the **-mfxsr**, **-mxsave**, and **-mxsaveopt** command line options respectively.

In addition, support for the **RDSEED**, **ADCX**, **ADOX**, and **PREFETCHW** instructions has been added and can be enabled by using the **-mrdseed**, **-madx**, and **-mprfchw** command line options.

New Built-in Functions to Detect Run-time CPU Type and ISA

A new built-in function, **__builtin_cpu_is()**, has been added to detect if the run-time CPU is of a particular type. This function accepts one string literal argument with the CPU name, and returns a positive integer on a match and zero otherwise. For example, **__builtin_cpu_is("westmere")** returns a positive integer if the run-time CPU is an Intel Core i7 Westmere processor. For a complete list of valid CPU names, see the user manual.

A new built-in function, **__builtin_cpu_supports()**, has been added to detect if the run-time CPU supports a particular ISA feature. This function accepts one string literal argument with the ISA feature, and returns a positive integer on a match and zero otherwise. For example, **__builtin_cpu_supports("ssse3")** returns a positive integer if the run-time CPU supports **SSSE3** instructions. For a complete list of valid ISA names, see the user manual.



Important

If these built-in functions are called before any static constructors are invoked, such as **IFUNC** initialization, then the CPU detection initialization must be explicitly run using this newly provided built-in function, **__builtin_cpu_init()**. The initialization needs to be done only once. For example, the following is sample invocation inside an **IFUNC** initializer:

```
static void (*some_ifunc_resolver(void))(void)
{
    __builtin_cpu_init();
    if (__builtin_cpu_is("amdfam10h") ...
    if (__builtin_cpu_supports("popcnt") ...
}
```

Function Multiversioning

Function multiversioning allows the programmer to specify multiple versions of the same function, each of which is specialized for a particular variant of a given target. At runtime, the appropriate version is automatically executed depending upon the target where the execution takes place. For example, consider the following code fragment:

```
__attribute__((target ("default"))) int foo () { return 0; }
__attribute__((target ("sse4.2"))) int foo () { return 1; }
__attribute__((target ("arch=atom"))) int foo () { return 2; }
```

When the function **foo()** is executed, the result returned depends upon the architecture where the program runs, not the architecture where the program was compiled. See the [GCC Wiki](#) for more details.

New RTM and HLE Intrinsics

Support for the Intel RTM and HLE intrinsics, built-in functions, and code generation has been added and can be enabled by using the **-mrtm** and **-mhle** command line options. This is done via intrinsics for *Restricted Transactional Memory* (RTM) and extensions to the memory model for *Hardware Lock Elision* (HLE).

For HLE, two new flags can be used to mark a lock as using hardware elision:

__ATOMIC_HLE_ACQUIRE

Starts lock elision on a lock variable. The memory model in use must be **__ATOMIC_ACQUIRE** or stronger.

__ATOMIC_HLE_RELEASE

Ends lock elision on a lock variable. The memory model must be **__ATOMIC_RELEASE** or stronger.

For example, consider the following code fragment:

```
while (__atomic_exchange_n (& lockvar, 1, __ATOMIC_ACQUIRE
                           | __ATOMIC_HLE_ACQUIRE))
    _mm_pause ();

// work with the acquired lock

__atomic_clear (& lockvar, __ATOMIC_RELEASE | __ATOMIC_HLE_RELEASE);
```

The new intrinsics that support Restricted Transactional Memory are:

unsigned _xbegin (void)

Attempts to start a transaction. If it succeeds, this function returns **_XBEGIN_STARTED**, otherwise it returns a status value indicating why the transaction could not be started.

void _xend (void)

Commits the current transaction. When no transaction is active, this function causes a fault. All memory side effects of the transactions become visible to other threads in an atomic manner.

int _xtest (void)

Returns a non-zero value if a transaction is currently active, or zero if it is not.

void _xabort (unsigned char status)

Aborts the current transaction. When no transaction is active, this is a no-op. The parameter **status** is included in the return value of any **_xbegin()** call that is aborted by this function.

The following example illustrates the use of these intrinsics:

```
if ((status = _xbegin ()) == _XBEGIN_STARTED)
{
    // some code
    _xend ();
}
else
{
    // examine the status to see why the transaction failed and possibly retry
}
```

Transactions Using Transactional Synchronization Extensions

Transactions in the transactional memory feature (the **-fgnu-tm** option) of GCC can now be run using *Transactional Synchronization Extensions* (TSX) if available on x86 hardware.

Support for AMD Family 15h Processors

The x86 backend of GCC now supports CPUs based on AMD Family 15h cores with the 64-bit x86 instruction set support. This can be enabled by using the **-march=bdver3** option.

Support for AMD Family 16h Processors

The x86 backend of GCC now supports CPUs based on AMD Family 16h cores with the 64-bit x86 instruction set support. This can be enabled by using the **-march=btver2** option.

B.2.2. Changes Since Red Hat Enterprise Linux 6.4 and 5.9

The following features have been added since the release of GCC included in Red Hat Enterprise Linux 6.4 and 5.9:

B.2.2.1. Status and Features

B.2.2.1.1. C++11

GCC 4.7 and later provides experimental support for building applications compliant with C++11 using the **-std=c++11** or **-std=gnu++11** command line options. However, there is no guarantee for compatibility between C++11 code compiled by different versions of the compiler. Refer to [Section B.2.2.3.1, “C++ ABI”](#) for details.

The C++ runtime library, **libstdc++**, supports a majority of the C++11 features. However, there is no or only partial support for some features such as certain properties on type traits or regular expressions. For details, refer to the [libstdc++ documentation](#), which also lists implementation-defined behavior.

Support for C++11 **exception_ptr** and **future** requires changes to the exception handling runtime in the system **libstdc++** package. These changes will be distributed through the normal Z-stream channel. Application of all Red Hat Enterprise Linux errata may be required to see correct runtime functionality

when using these features.

B.2.2.1.2. C11

GCC 4.7 and later provides experimental support for some of the features from the C11 revision of the ISO C standard, and in addition to the previous (now deprecated) **-std=c1x** and **-std=gnu1x** command line options, gcc now accepts **-std=c11** and **-std=gnu11**. Note that since this support is experimental, it may change incompatibly in future releases.

Examples for features that are supported are Unicode strings (including the predefined macros `__STDC_UTF_16__` and `__STDC_UTF_32__`), nonreturning functions (`_Noreturn` and `<stdnoreturn.h>`), and alignment support (`_Alignas`, `_Alignof`, `max_align_t`, and `<stdalign.h>`).

B.2.2.1.3. Parallelism and Concurrency

GCC 4.7 and later provides improved support for programming parallel applications:

1. The GCC compilers support the OpenMP API specification for parallel programming, version 3.1. Refer to the [OpenMP](#) website for more information about this specification.
2. The C++11 and C11 standards provide programming abstractions for multi-threaded programs. The respective standard libraries include programming abstractions for threads and thread-related features such as locks, condition variables, or futures. These new versions of the standard also define a memory model that precisely specifies the runtime behavior of a multi-threaded program, such as the guarantees provided by compilers and the constraints programmers have to pay attention to when writing multi-threaded programs.
Note that support for the memory model is still experimental (see below for details). For more information about the status of support for C++11 and C11, refer to [Section B.2.2.1.1, “C++11”](#) and [Section B.2.2.1.2, “C11”](#) respectively.

The rest of this section describes two new GCC features in more detail. Both these features make it easier for programmers to handle concurrency (such as when multiple threads do not run truly in parallel but instead have to synchronize concurrent access to shared state), and both provide atomicity for access to memory but differ in their scope, applicability, and complexity of runtime support.

C++11 Types and GCC Built-ins for Atomic Memory Access

C++11 has support for *atomic types*. Access to memory locations of this type is atomic, and appears as one indivisible access even when other threads access the same memory location concurrently. The atomicity is limited to a single read or write access or one of the other atomic operations supported by such types (for example, two subsequent operations executed on a variable of atomic type are each atomic separately, but do not form one joint atomic operation).

An atomic type is declared as **atomic<T>**, where **T** is the non-atomic base type and must be trivially copyable (for example, **atomic<int>** is an atomic integer). GCC does not yet support any base type **T**, but only those that can be accessed atomically with the atomic instructions offered by the target architecture. This is not a significant limitation in practice, given that atomics are primarily designed to expose hardware primitives in an architecture-independent fashion; pointers and integrals that are not larger than a machine word on the target are supported as base types. Using base types that are not yet supported results in link-time errors.

The code generated for operations on atomic types, including the memory orders, implements the semantics specified in the C++11 standard. However, support for the C++11 memory model is still experimental, and for example GCC might not always preserve data-race freedom when optimizing code.

GCC also supports new built-ins for atomic memory accesses, which follow the design of the memory model and new atomic operations. The former set of synchronization built-ins (that is, those prefixed

with `__sync`) are still supported.

Transactional Memory

Transactional Memory (TM) allows programs to declare that a piece of code is supposed to execute as a transaction, that is, virtually atomically and in isolation from other transactions. GCC's transactional memory runtime library, `libitm`, then ensures this atomicity guarantee when executing the compiled program. Compared to atomic memory accesses, it is a higher-level programming abstraction, because it is not limited to single memory locations, does not require special data types for the data it modifies, and because transactions can contain arbitrary code and be nested within other transactions (with some restrictions explained subsequently).

GCC implements transactions as specified in the [Draft Specification for Transactional Language Constructs for C++, version 1.1](#). This draft does not yet specify the language constructs for C, but GCC already supports a C-compatible subset of the constructs when compiling C source code.

The main language constructs are transaction statements and expressions, and are declared by the `__transaction_atomic` or `__transaction_relaxed` keywords followed by a compound statement or expression, respectively. The following example illustrates how to increment a global variable `y` if another variable `x` has a value less than 10:

```
__transaction_atomic { if (x < 10) y++; }
```

This happens atomically even in a multi-threaded execution of the program. In particular, even though the transaction can load `x` and `y` and store to `y`, all these memory accesses are virtually executed as one indivisible step.

Note that in line with the C++11 memory model, programs that use transactions must be free of data races. Transactions are guaranteed to be virtually executed serially in a global total order that is determined by the transactional memory implementation and that is consistent with and contributes to the happens-before order enforced by the rest of the program (that is, transaction semantics are specified based on the C++11 memory model, see the draft specification linked above). Nonetheless, if a program is not data-race-free, then it has undefined behavior. For example, a thread can first initialize some data and then make it publicly accessible by code like this:

```
init(data);
__transaction_atomic { data_public = true; } // data_public is initially false
```

Another thread can then safely use the data, for instance:

```
__transaction_atomic { if (data_public) use(data); }
```

However, the following code has a data race and thus results in undefined behavior:

```
__transaction_atomic { temp = copy(data); if (data_public) use(temp); }
```

Here, `copy(data)` races with `init(data)` in the initializing thread, because this can be executed even if `data_public` is not true. Another example for data races is one thread accessing a variable `x` transactionally and another thread accessing it nontransactionally at potentially the same time. Note that the data can be safely reclaimed using code like this (assuming only one thread ever does this):

```
__transaction_atomic { data_public = false; }
destruct(data);
```

Here, **destruct()** does not race with potential concurrent uses of the data because after the transaction finishes, it is guaranteed that **data_public** is false and thus data is private. See the specification and the C++11 memory model for more background information about this.

Note that even if transactions are required to virtually execute in a total order, this does not mean that they execute mutually exclusive in time. Transactional memory implementations attempt to run transactions as much in parallel as possible to provide scalable performance.

There are two variants of transactions: *atomic transactions* (**__transaction_atomic**) and *relaxed transactions* (**__transaction_relaxed**). The former guarantee atomicity with regard to all other code, but allow only code that is known to not include nontransactional kinds of synchronization, such as atomic or volatile memory access. In contrast, relaxed transactions allow all code (for example calls to I/O functions), but only provide atomicity with regard to other transactions. Therefore, atomic transactions can be nested within other atomic and relaxed transactions, but relaxed transactions can only be nested within other relaxed transactions. Furthermore, relaxed transactions are likely to be executed with less performance, but this depends on the implementation and available hardware.

GCC verifies these restrictions statically at compile time (for example, the requirements on code allowed to be called from within atomic transactions). This has implications for when transactions call functions that are defined within other compilation unit (source file) or within libraries. To enable such cross-compilation-unit calls for transactional code, the respective functions must be marked to contain code that is safe to use from within atomic transactions. Programmers can do so by adding the **transaction_safe** function attribute to the declarations of these functions and by including this declaration when defining the function. In turn, GCC then verifies that the code in these functions is safe for atomic transactions and generates code accordingly. If the programmer does not follow these constraints and/or steps, compile-time or link-time errors occur. Note that within a compilation unit, GCC detects automatically whether a function is safe for use within transactions, and the attributes therefore typically do not need to be added. See the draft specification linked above for further details.

GCC's transactional memory support is designed in such a way that it does not decrease the performance of programs that do not use transactions, nor the performance of nontransactional code, except due to the normal kinds of interference by concurrent threads that use the same resources such as the CPU.

Transactional memory support in GCC and **libitm** is still experimental, and both the ABI and API could change in the future if this is required due to the evolution of the specification of the language constructs, or due to implementation requirements. Note that when executing applications built with the **-fgnu-tm** command line option, it is currently a prerequisite to also have the appropriate version of the **libitm.so.1** shared library installed.

B.2.2.1.4. Architecture-specific Options

Red Hat Developer Toolset 2.0 is only available for Red Hat Enterprise Linux 5 and 6, both for the 32-bit and 64-bit Intel and AMD architectures. Consequently, the options described below are only relevant to these architectures.

Optimization for several processors is now available through the command line options described in [Table B.1, “Processor Optimization Options”](#).

Table B.1. Processor Optimization Options

| Option | Description |
|---|---|
| -march=core2 and -mtune=core2 | Optimization for Intel Core 2 processors. |
| -march=corei7 and -mtune=corei7 | Optimization for Intel Core i3, i5, and i7 processors. |
| -march=corei7-avx and -mtune=corei7-avx | Optimization for Intel Core i3, i5, and i7 processors with AVX. |
| -march=core-avx-i | Optimization for the Intel processor code-named Ivy Bridge with RDRND, FSGSBASE, and F16C. |
| -march=core-avx2 | Optimization for a next-generation processor from Intel with AVX2, FMA, BMI, BMI2, and LZCNT. |
| -march=bdver2 and -mtune=bdver2 | Optimization for AMD Opteron processors code-named Piledriver. |
| -march=btver1 and -mtune=btver1 | Optimization for AMD family 14 processors code-named Bobcat. |
| -march=bdver1 and -mtune=bdver1 | Optimization for AMD family 15h processors code-named Bulldozer. |

Support for various processor-specific intrinsics and instructions is now available through the command line options described in [Table B.2, “Support for Processor-specific Intrinsics and Instructions”](#).

Table B.2. Support for Processor-specific Intrinsics and Instructions

| Option | Description |
|---|--|
| -mavx2 | Support for Intel AVX2 intrinsics, built-in functions, and code generation. |
| -mbmi2 | Support for Intel BMI2 intrinsics, built-in functions, and code generation. |
| -mlzcnt | Implementation and automatic generation of __builtin_clz* using the lzcnt instruction. |
| -mfma | Support for Intel FMA3 intrinsics and code generation. |
| -mfsgsbase | Enables the generation of new segment register read/write instructions through dedicated built-ins. |
| -mrdrnd | Support for the Intel rdrnd instruction. |
| -mf16c | Support for two additional AVX vector conversion instructions. |
| -mtbm | Support for TBM (Trailing Bit Manipulation) built-in functions and code generation. |
| -mbmi | Support for AMD's BMI (Bit Manipulation) built-in functions and code generation. |
| -mcrc32 | Support for crc32 intrinsics. |
| -mmovbe | Enables the use of the movbe instruction to implement __builtin_bswap32 and __builtin_bswap64 . |
| -mxop , -mfma4 , and -mlwp | Support for the XOP, FMA4, and LWP instruction sets for the AMD Orochi processors. |
| -mabm | Enables the use of the popcnt and lzcnt instructions on AMD processors. |
| -mpopcnt | Enables the use of the popcnt instruction on both AMD and Intel processors. |

When using the x87 floating-point unit, GCC now generates code that conforms to ISO C99 in terms of handling of floating-point excess precision. This can be enabled by **-fexcess-precision=standard** and disabled by **-fexcess-precision=fast**. This feature is enabled by default when using standards conformance options such as **-std=c99**.

Vectors of type **vector long long** or **vector long** are passed and returned using the same method as other vectors with the **VSX** instruction set. Previously GCC did not adhere to the ABI for 128-bit vectors with 64-bit integer base types (see GCC PR 48857).

The **-mrecip** command line option has been added, which indicates whether the reciprocal and reciprocal square root instructions should be used.

The **-mveclibabi=mass** command line option has been added. This can be used to enable the compiler to auto-vectorize mathematical functions using the Mathematical Acceleration Subsystem library.

The **-msingle-pic-base** command line option has been added, which instructs the compiler to avoid loading the **PIC** base register in function prologues. The PIC base register must be initialized by the runtime system.

The **-mblock-move-inline-limit** command line option has been added, which enables the user to control the maximum size of inlined **memcpy** calls and similar.

B.2.2.1.5. Link-time Optimization

Link-time optimization (LTO) is a compilation technique in which GCC generates an internal representation of each compiled input file in addition to the native code, and writes both to the output object file. Subsequently, when several object files are linked together, GCC uses the internal representations of the compiled code to optimize inter-procedurally across all the compilation units. This can potentially improve the performance of the generated code (for example, functions defined in one file can potentially be inlined when called in another file).

To enable LTO, the **-flto** option needs to be specified at both compile time and link time. For further details, including interoperability with linkers and parallel execution of LTO, refer to the documentation for **-flto** in the [GCC 4.7.0 Manual](#). Also note that the internal representation is not a stable interface, so LTO will only apply to code generated by the same version of GCC.



Note

Use of Link-time Optimization with debug generation is not yet supported in gcc 4.7 and 4.8 and so use of the **-flto** and the **-g** options together is unsupported in Red Hat Developer Toolset.

B.2.2.1.6. Miscellaneous

-Ofast is now supported as a general optimization level. It operates similar to **-O3**, adds options that can yield better-optimized code, but in turn might invalidate standards compliance (for example, **-ffast-math** is enabled by **-Ofast**).

GCC can now inform users about cases in which code generation might be improved by adding attributes such as **const**, **pure**, and **noreturn** to functions declared in header files. Use the **-Wsuggest-attribute=[const|pure|noreturn]** command line option to enable this.

Assembler code can now make use of a **goto** feature that allows for jumps to labels in C code.

B.2.2.2. Language Compatibility

In this section, we describe the compatibility between the Red Hat Developer Toolset compilers and the Red Hat Enterprise Linux system compilers at the programming-language level (for example, differences in the implementation of language standards such as C99, or changes to the warnings generated by `-Wall`).

Some of the changes are a result of bug fixing, and some old behaviors have been intentionally changed in order to support new standards, or relaxed in standards-conforming ways to facilitate compilation or runtime performance. Some of these changes are not visible to the naked eye and will not cause problems when updating from older versions. However, some of these changes are visible, and can cause grief to users porting to Red Hat Developer Toolset's version of GCC. The following text attempts to identify major issues and suggests solutions.

B.2.2.2.1. C

Constant expressions are now handled by GCC in a way that conforms to C90 and C99. For code expressions that can be transformed into constants by the compiler but are in fact not constant expressions as defined by ISO C, this may cause warnings or errors.

Ill-formed redeclarations of library functions are no longer accepted by the compiler. In particular, a function with a signature similar to the built-in declaration of a library function (for example, `abort()` or `memcpy()`) must be declared with `extern "C"` to be considered as a redeclaration, otherwise it is ill-formed.

Duplicate Member

Consider the following `struct` declaration:

```
struct A { int *a; union { struct { int *a; }; }; };
```

Previously, this declaration used to be diagnosed just by the C++ compiler, now it is also diagnosed by the C compiler. Because of the anonymous unions and structs, there is ambiguity about what `.a` actually refers to and one of the fields therefore needs to be renamed.

B.2.2.2.2. C++

Header Dependency Changes

`<iostream>`, `<string>`, and other STL headers that previously included `<unistd.h>` as an implementation detail (to get some feature macros for `gthr*.h` purposes) no longer do so, because it was a C++ standard violation. This can result in diagnostic output similar to the following:

```
error: 'truncate' was not declared in this scope
error: 'sleep' was not declared in this scope
error: 'pipe' was not declared in this scope
error: there are no arguments to 'offsetof' that depend on a template
parameter, so a declaration of 'offsetof' must be available
```

To fix this, add the following line early in the source or header files that need it:

```
#include <unistd.h>
```

Many of the standard C++ library include files have been edited to no longer include `<cstdint>` to get namespace-`std`-scoped versions of `size_t` and `ptrdiff_t`. As such, C++ programs that used the

macros **NULL** or **offsetof** without including **<cstddef>** will no longer compile. The diagnostic produced is similar to the following:

```
error: 'ptrdiff_t' does not name a type
error: 'size_t' has not been declared
error: 'NULL' was not declared in this scope
error: there are no arguments to 'offsetof' that depend on a template
parameter, so a declaration of 'offsetof' must be available
```

To fix this issue, add the following line:

```
#include <cstddef>
```

Name Lookup Changes

G++ no longer performs an extra unqualified lookup that it incorrectly performed in the past. Instead, it implements the two-phase lookup rules correctly, and an unqualified name used in a template must have an appropriate declaration that:

1. is either in scope at the point of the template's definition, or
2. can be found by argument-dependent lookup at the point of instantiation.

Code that incorrectly depends on a second unqualified lookup at the point of instantiation (such as finding functions declared after the template or in dependent bases) will result in compile-time errors.

In some cases, the diagnostics provided by G++ include hints how to fix the bugs. Consider the following code:

```
template<typename T>
int t(T i)
{
    return f(i);
}

int f(int i)
{
    return i;
}

int main()
{
    return t(1);
}
```

The following diagnostics output will be produced:

```
In instantiation of 'int t(T) [with T = int]'
required from here
error: 'f' was not declared in this scope, and no declarations were found by
argument-dependent lookup at the point of instantiation [-fpermissive]
note: 'int f(int)' declared here, later in the translation unit
```

To correct the error in this example, move the declaration of function **f()** before the definition of template function **t()**. The **-fpermissive** compiler flag turns compile-time errors into warnings and can be used as a temporary workaround.

Uninitialized const

Consider the following declaration:

```
struct A { int a; A (); };
struct B : public A { };
const B b;
```

An attempt to compile this code now fails with the following error:

```
error: uninitialized const 'b' [-fpermissive]
note: 'const struct B' has no user-provided default constructor
```

This happens, because **B** does not have a user-provided default constructor. Either an initializer needs to be provided, or the default constructor needs to be added.

Visibility of Template Instantiations

The ELF symbol visibility of a template instantiation is now properly constrained by the visibility of its template arguments. For instance, users that instantiate standard library components like `std::vector` with hidden user defined types such as `struct my_hidden_struct` can now expect hidden visibility for `std::vector<my_hidden_struct>` symbols. As a result, users that compile with the `-fvisibility=hidden` command line option should be aware of the visibility of types included from the library headers used. If the header does not explicitly control symbol visibility, types from those headers will be hidden, along with instantiations that use those types. For instance, consider the following code:

```
#include <vector> // template std::vector has default visibility
#include <ctime> // struct tm has hidden visibility
template class std::vector<tm>; // instantiation has hidden visibility
```

One approach to adjusting the visibility of a library header `<foo.h>` is to create a forwarding header on the `-I` include path consisting of the following:

```
#pragma GCC visibility push(default)
#include_next <foo.h>
#pragma GCC visibility push
```

User-defined Literal Support

When compiling C++ with the `-std={c++11,c++0x,gnu++11,gnu++0x}` command line option, GCC 4.7.0 and later, unlike older versions, supports user-defined literals, which are incompatible with some valid ISO C++03 code. In particular, white space is now needed after a string literal before something that could be a valid user defined literal. Consider the following code:

```
const char *p = "foobar"__TIME__;
```

In C++03, the `__TIME__` macro expands to some string literal and is concatenated with the other one. In C++11, `__TIME__` is not expanded and instead, operator `" " __TIME__` is being looked up, which results in a warning like:

```
error: unable to find string literal operator 'operator"" __TIME__'
```

This applies to any string literal followed without white space by some macro. To fix this, add some white space between the string literal and the macro name.

Taking the Address of Temporary

Consider the following code:

```
struct S { S (); int i; };
void bar (S *);
void foo () { bar (&S ()); }
```

Previously, an attempt to compile this code produced a warning message, now it fails with an error. This can be fixed by adding a variable and passing the address of this variable instead of the temporary. The **-fpermissive** compiler flag turns compile-time errors into warnings and can be used as a temporary workaround.

Miscellaneous

G++ now sets the predefined macro **__cplusplus** to the correct value: **199711L** for C++98/03, and **201103L** for C++11.

G++ now properly re-uses stack space allocated for temporary objects when their lifetime ends, which can significantly lower stack consumption for some C++ functions. As a result of this, some code with undefined behavior will now break.

When an extern declaration within a function does not match a declaration in the enclosing context, G++ now properly declares the name within the namespace of the function rather than the namespace which was open just before the function definition.

G++ now implements the proposed resolution of the C++ standard's core issue 253. Default initialization is allowed if it initializes all subobjects, and code that fails to compile can be fixed by providing an initializer such as:

```
struct A { A(); };
struct B : A { int i; };
const B b = B();
```

Access control is now applied to **typedef** names used in a template, which may cause G++ to reject some ill-formed code that was accepted by earlier releases. The **-fno-access-control** option can be used as a temporary workaround until the code is corrected.

G++ now implements the C++ standard's core issue 176. Previously, G++ did not support using the injected-class-name of a template base class as a type name, and lookup of the name found the declaration of the template in the enclosing scope. Now lookup of the name finds the injected-class-name, which can be used either as a type or as a template, depending on whether or not the name is followed by a template argument list. As a result of this change, some code that was previously accepted may be ill-formed, because:

1. the injected-class-name is not accessible because it is from a private base, or
2. the injected-class-name cannot be used as an argument for a template parameter.

In either of these cases, the code can be fixed by adding a nested-name-specifier to explicitly name the template. The first can be worked around with **-fno-access-control**, the second is only rejected with **-pedantic**.

B.2.2.2.3. C/C++ Warnings

GCC 4.7.0 and later adds a number of new warnings that are either enabled by default, or by using the **-Wall** option. Although these warnings do not result in a compilation failure on their own, often **-Wall** is

used in conjunction with **-Werror**, causing these warnings to act like errors. This section provides a list of these new or newly enabled warnings. Unless noted otherwise, these warnings apply to both C and C++.

The behavior of the **-Wall** command line option has changed and now includes the new warning flags **-Wunused-but-set-variable** and, with **-Wall -Wextra**, **-Wunused-but-set-parameter**. This may result in new warnings in code that compiled cleanly with previous versions of GCC. For example, consider the following code:

```
void fn (void)
{
    int foo;
    foo = bar (); /* foo is never used. */
}
```

The following diagnostic output will be produced:

```
warning: variable "foo" set but not used [-Wunused-but-set-variable]
```

To fix this issue, first see if the unused variable or parameter can be removed without changing the result or logic of the surrounding code. If not, annotate it with `__attribute__((__unused__))`. As a workaround, you can use the **-Wno-error=unused-but-set-variable** or **-Wno-error=unused-but-set-parameter** command line option.

The **-Wenum-compare** option causes GCC to report a warning when values of different enum types are being compared. Previously, this option only worked for C++ programs, but now it works for C as well. This warning is enabled by **-Wall** and may be avoided by using a type cast.

Casting integers to larger pointer types now causes GCC to display a warning by default. To disable these warnings, use the **-Wno-int-to-pointer-cast** option, which is available for both C and C++.

Conversions between NULL and non-pointer types now cause GCC to report a warning by default. Previously, these warnings were only displayed when explicitly using **-Wconversion**. To disable these warnings, use the new **-Wno-conversion-null** command line option.

GCC can now warn when a class that has virtual functions and a non-virtual destructor is destroyed by using **delete**. This is unsafe to do because the pointer might refer to a base class that does not have a virtual destructor. The warning is enabled by **-Wall** and by a new command line option, **-Wdelete-non-virtual-dtor**.

New **-Wc++11-compat** and **-Wc++0x-compat** options are now available. These options cause GCC to display a warning about C++ constructs whose meaning differs between ISO C++ 1998 and ISO C++ 2011 (such as identifiers in ISO C++ 1998 that are keywords in ISO C++ 2011). This warning is enabled by **-Wall** and enables the **-Wnarrowing** option.

B.2.2.2.4. Fortran

B.2.2.2.4.1. New Features

- ▶ A new compile flag **-fstack-arrays** has been added. This flag causes all local arrays to be put on stack memory, which can significantly improve the performance of some programs. Note that programs that use very large local arrays may require you to extend your runtime limits for stack memory.
- ▶ Compile time has been significantly improved. For example, the improvement may be noticeable when

working with programs that use large array constructors.

- ▶ To improve code generation and diagnostics, the **-fwhole-file** compile flag is now enabled by default, and can be used with a newly supported **-fwhole-program** flag. To disable it, use the deprecated **-fno-whole-file** flag.
- ▶ A new command line option **-M** is now supported. Similarly to **gcc**, this option allows you to generate Makefile dependencies. Note that the **-cpp** option may be required as well.
- ▶ The **-finit-real=** command line option now supports **snan** as a valid value. This allows you to initialize REAL and COMPLEX variables with a signaling NaN (*not a number*), and requires you to enable trapping (for example, by using the **-ffpe-trap=** command line option). Note that compile-time optimizations may turn a signaling NaN into a quiet NaN.
- ▶ A new command line option **-fcheck=** has been added. This option accepts the following arguments:
 - The **-fcheck=bounds** option is equivalent to the **-fbounds-check** command line option.
 - The **-fcheck=array-temps** option is equivalent to the **-fcheck-array-temporaries** command line option.
 - The **-fcheck=do** option checks for invalid modification of loop iteration variables.
 - The **-fcheck=recursive** option checks for recursive calls to subroutines or functions that are not marked as recursive.
 - The **-fcheck=pointer** option performs pointer association checks in calls, but does not handle undefined pointers nor pointers in expressions.
 - The **-fcheck=all** option enables all of the above options.
- ▶ A new command line option **-fno-protect-parens** has been added. This option allows the compiler to reorder REAL and COMPLEX expressions with no regard to parentheses.
- ▶ When OpenMP's **WORKSHARE** is used, array assignments and **WHERE** will now be run in parallel.
- ▶ More Fortran 2003 and Fortran 2008 mathematical functions can now be used as initialization expressions.
- ▶ The **GCC\$** compiler directive now enables support for some extended attributes such as **STDCALL**.

B.2.2.2.4.2. Compatibility Changes

- ▶ The **-Ofast** command line option now automatically enables the **-fno-protect-parens** and **-fstack-arrays** flags.
- ▶ Front-end optimizations can now be disabled by the **-fno-frontend-optimize** option, and selected by the **-ffrontend-optimize** option. The former is essentially only desirable if invalid Fortran source code needs to be compiled (for example, when functions—as compared to subroutines—have side-effects) or to work around compiler bugs.
- ▶ The **GFORTRAN_USE_STDERR** environment variable has been removed, and GNU Fortran now always prints error messages to standard error.
- ▶ The **-fdump-core** command line option and the **GFORTRAN_ERROR_DUMP CORE** environment variable have been removed. When encountering a serious error, GNU Fortran now always aborts the execution of the program.
- ▶ The **-fbacktrace** command line option is now enabled by default. When a fatal error occurs, GNU Fortran now attempts to print a backtrace to standard error before aborting the execution of the program. To disable this behavior, use the **-fno-backtrace** option.
- ▶ GNU Fortran no longer supports the use of the **-M** command line option to generate Makefile dependencies for the module path. To perform this operation, use the **-J** option instead.
- ▶ To significantly reduce the number of warnings, the **-Wconversion** command line option now only

displays warnings when a conversion leads to information loss, and a new command line option **-Wconversion-extra** has been added to display warnings about other conversions. The **-Wconversion** option is now enabled with **-Wall**.

- ▶ A new command line option **-Wunused-dummy-argument** has been added. This option can be used to display warnings about unused dummy arguments, and is now enabled with **-Wall**. Note that the **-Wunused-variable** option previously also warned about unused dummy arguments.
- ▶ The **COMMON** default padding has been changed. Previously, the padding was added before a variable. Now it is added after a variable to increase the compatibility with other vendors, as well as to help to obtain the correct output in some cases. Note that this behavior is in contrast with the behavior of the **-falign-commons** option.
- ▶ GNU Fortran no longer links against the **libgfortranbegin** library. The **MAIN__** assembler symbol is the actual Fortran main program and is invoked by the **main** function, which is now generated and put in the same object file as **MAIN__**. Note that the **libgfortranbegin** library is still present for backward compatibility.

B.2.2.2.4.3. Fortran 2003 Features

- ▶ Improved but still experimental support for polymorphism between libraries and programs and for complicated inheritance patterns.
- ▶ Generic interface names which have the same name as derived types are now supported, which allows the creation of constructor functions. Note that Fortran does not support static constructor functions; only default initialization or an explicit structure-constructor initialization are available.
- ▶ Automatic (re)allocation: In intrinsic assignments to allocatable variables, the left-hand side will be automatically allocated (if unallocated) or reallocated (if the shape or type parameter is different). To avoid the small performance penalty, you can use **a(:) = ...** instead of **a = ...** for arrays and character strings — or disable the feature using **-std=f95** or **-fno-realloc-lhs**.
- ▶ Experimental support of the **ASSOCIATE** construct has been added.
- ▶ In pointer assignments it is now possible to specify the lower bounds of the pointer and, for a rank-1 or a simply contiguous data-target, to remap the bounds.
- ▶ Deferred type parameter: For scalar allocatable and pointer variables the character length can now be deferred.
- ▶ Namelist variables with allocatable attribute, pointer attribute, and with a non-constant length type parameter are now supported.
- ▶ Support has been added for procedure-pointer function results and procedure-pointer components (including **PASS**).
- ▶ Support has been added for allocatable scalars (experimental), **DEFERRED** type-bound procedures, and the **ERRMSG=** argument of the **ALLOCATE** and **DEALLOCATE** statements.
- ▶ The **ALLOCATE** statement now supports type-specs and the **SOURCE=** argument.
- ▶ Rounding (**ROUND=**, **RZ**, ...) for output is now supported.
- ▶ The **INT_FAST{8,16,32,64,128}_T** format for **ISO_C_BINDING** intrinsic module type parameters is now supported.
- ▶ **OPERATOR(*)** and **ASSIGNMENT(=)** are now allowed as **GENERIC** type-bound procedures (i.e. as type-bound operators).

B.2.2.2.4.4. Fortran 2003 Compatibility

Extensible derived types with type-bound procedure or procedure pointer with **PASS** attribute now have to use **CLASS** in line with the Fortran 2003 standard; the workaround to use **TYPE** is no longer supported.

B.2.2.2.4.5. Fortran 2008 Features

- ▶ A new command line option **-std=f2008ts** has been added. This option enables support for programs that conform to the Fortran 2008 standard and the draft Technical Specification (TS) 29113 on Further Interoperability of Fortran with C. For more information, refer to the [Chart of Fortran TS 29113 Features supported by GNU Fortran](#).
- ▶ The **DO CONCURRENT** construct is now supported. This construct can be used to specify that individual loop iterations do not have any interdependencies.
- ▶ Full single-image support except for polymorphic coarrays has been added, and can be enabled by using the **-fcoarray=single** command line option. Additionally, GNU Fortran now provides preliminary support for multiple images via an MPI-based coarray communication library. Note that the library version is not yet usable as remote coarray access is not yet possible.
- ▶ The **STOP** and **ERROR STOP** statements have been updated to support all constant expressions.
- ▶ The **CONTIGUOUS** attribute is now supported.
- ▶ Use of **ALLOCATE** with the **MOLD** argument is now supported.
- ▶ The **STORAGE_SIZE** intrinsic inquiry function is now supported.
- ▶ The **NORM2** and **PARITY** intrinsic functions are now supported.
- ▶ The following bit intrinsics have been added:
 - the **POPCNT** and **POPPAR** bit intrinsics for counting the number of 1 bits and returning the parity;
 - the **BGE**, **BGT**, **BLE**, and **BLT** bit intrinsics for bitwise comparisons;
 - the **DSHIFTL** and **DSHIFTR** bit intrinsics for combined left and right shifts;
 - the **MASKL** and **MASKR** bit intrinsics for simple left and right justified masks;
 - the **MERGE_BITS** bit intrinsic for a bitwise merge using a mask;
 - the **SHIFTA**, **SHIFTL**, and **SHIFTR** bit intrinsics for shift operations;
 - the transformational bit intrinsics **IALL**, **IANY**, and **IPARITY**.
- ▶ The **EXECUTE_COMMAND_LINE** intrinsic subroutine is now supported.
- ▶ The **IMPURE** attribute for procedures is now supported. This allows the use of **ELEMENTAL** procedures without the restrictions of **PURE**.
- ▶ Null pointers (including **NULL()**) and unallocated variables can now be used as an actual argument to optional non-pointer, non-allocatable dummy arguments, denoting an absent argument.
- ▶ Non-pointer variables with the **TARGET** attribute can now be used as an actual argument to **POINTER** dummies with **INTENT(IN)**.
- ▶ Pointers that include procedure pointers and those in a derived type (pointer components) can now also be initialized by a target instead of only by **NULL**.
- ▶ The **EXIT** statement (with construct-name) can now be used to leave the **ASSOCIATE**, **BLOCK**, **IF**, **SELECT CASE**, and **SELECT TYPE** constructs in addition to **DO**.
- ▶ Internal procedures can now be used as actual arguments.
- ▶ The named constants **INTEGER_KINDS**, **LOGICAL_KINDS**, **REAL_KINDS**, and **CHARACTER_KINDS** of the intrinsic module **ISO_FORTRAN_ENV** have been added. These arrays contain the supported 'kind' values for the respective types.
- ▶ The **C_SIZEOF** module procedures of the **ISO_C_BINDINGS** intrinsic module and the **COMPILER_VERSION** and **COMPILER_OPTIONS** module procedures of the **ISO_FORTRAN_ENV** intrinsic module have been implemented.
- ▶ The **OPEN** statement now supports the **NEWUNIT=** option. This option returns a unique file unit and therefore prevents inadvertent use of the same unit in different parts of the program.
- ▶ Unlimited format items are now supported.

- The **INT{8,16,32}** and **REAL{32,64,128}** format for **ISO_FORTRAN_ENV** intrinsic module type parameters are now supported.
- It is now possible to use complex arguments with the **TAN**, **SINH**, **COSH**, **TANH**, **ASIN**, **ACOS**, and **ATAN** functions. Additionally, the new functions **ASINH**, **ACOSH**, and **ATANH** have been added for real and complex arguments, and **ATAN(Y,X)** now serves as an alias for **ATAN2(Y,X)**.
- The **BLOCK** construct has been implemented.

B.2.2.2.4.6. Fortran 2008 Compatibility

The implementation of the **ASYNCHRONOUS** attribute in GCC is now compatible with the candidate draft of *TS 29113: Technical Specification on Further Interoperability with C*.

B.2.2.2.4.7. Fortran 77 Compatibility

When the GNU Fortran compiler is issued with the **-fno-sign-zero** option, the **SIGN** intrinsic now behaves as if zero were always positive.

B.2.2.3. ABI Compatibility

This section describes compatibility between the Red Hat Developer Toolset compilers and the system compilers at the *application binary interface* (ABI) level.

B.2.2.3.1. C++ ABI

Because the upstream GCC community development does not guarantee C++11 ABI compatibility across major versions of GCC, the same applies to use of C++11 with Red Hat Developer Toolset. Consequently, using the **-std=c++11** option is supported in Red Hat Developer Toolset 2.0 only when all C++ objects compiled with that flag have been built using the same major version of Red Hat Developer Toolset. The mixing of objects, binaries and libraries, built by the Red Hat Enterprise Linux 5 or 6 system toolchain GCC using the **-std=c++0x** or **-std=gnu++0x** flags, with those built with the **-std=c++11** or **-std=gnu++11** flags using the GCC in Red Hat Developer Toolset is explicitly not supported.

As later major versions of Red Hat Developer Toolset may use a later major release of GCC, forward-compatibility of objects, binaries, and libraries built with the **-std=c++11** or **-std=gnu++11** options cannot be guaranteed, and so is not supported.

The default language standard setting for Red Hat Developer Toolset is C++98. Any C++98-compliant binaries or libraries built in this default mode (or explicitly with **-std=c++98**) can be freely mixed with binaries and shared libraries built by the Red Hat Enterprise Linux 5 or 6 system toolchain GCC. Red Hat recommends use of this default **-std=c++98** mode for production software development.



Important

Use of C++11 features in your application requires careful consideration of the above ABI compatibility information.

Aside from the C++11 ABI, discussed above, [the Red Hat Enterprise Linux Application Compatibility Specification](#) is unchanged for Red Hat Developer Toolset. When mixing objects built with Red Hat Developer Toolset with those built with the Red Hat Enterprise Linux v5.x/v6.x toolchain (particularly .o/.a files), the Red Hat Developer Toolset toolchain should be used for any linkage. This ensures any newer library features provided only by Red Hat Developer Toolset are resolved at link-time.

A new standard mangling for SIMD vector types has been added to avoid name clashes on systems with

vectors of varying length. By default the compiler still uses the old mangling, but emits aliases with the new mangling on targets that support strong aliases. **-Wabi** will now display a warning about code that uses the old mangling.

B.2.2.3.2. Miscellaneous

GCC now optimizes calls to various standard C string functions such as **strlen()**, **strchr()**, **strcpy()**, **strcat()** and **strncpy()** (as well as their respective **_FORTIFY_SOURCE** variants) by transforming them into custom, faster code. This means that there might be fewer or other calls to those functions than in the original source code. The optimization is enabled by default at **-O2** or higher optimization levels. It is disabled when using **-fno-optimize-strlen** or when optimizing for size.

When compiling for 32-bit GNU/Linux and not optimizing for size, **-fomit-frame-pointer** is now enabled by default. The prior default setting can be chosen by using the **-fno-omit-frame-pointer** command line option.

Floating-point calculations on x86 targets and in strict C99 mode are now compiled by GCC with a stricter standard conformance. This might result in those calculations executing significantly slower. It can be disabled using **-fexcess-precision=fast**.

B.2.2.4. Debugging Compatibility

GCC now generates DWARF debugging information that uses more or newer DWARF features than previously. GDB contained in Red Hat Developer Toolset can handle these features, but versions of GDB older than 7.0 cannot. GCC can be restricted to only generate debugging information with older DWARF features by using the **-gdwarf-2 -gstrict-dwarf** or **-gdwarf-3 -gstrict-dwarf** options (the latter are handled partially by versions of GDB older than 7.0).

Many tools such as **Valgrind**, **SystemTap**, or third-party debuggers utilize debugging information. It is suggested to use the **-gdwarf-2 -gstrict-dwarf** options with those tools.



Note

Use of Link-time Optimization with debug generation is not yet supported in gcc 4.7 and 4.8 and so use of the **-flto** and the **-g** options together is unsupported in Red Hat Developer Toolset.

B.2.2.5. Other Compatibility

GCC is now more strict when parsing command line options, and both **gcc** and **g++** report an error when invalid command line options are used. In particular, when only linking and not compiling code, earlier versions of GCC ignored all options starting with **--**. For example, options accepted by the linker such as **--as-needed** and **--export-dynamic** are not accepted by **gcc** and **g++** anymore, and should now be directed to the linker using **-Wl, --as-needed** or **-Wl, --export-dynamic** if that is intended.

Because of the new link-time optimization feature (see [Section B.2.2.1.5, “Link-time Optimization”](#)), support for the older intermodule optimization framework has been removed and the **-combine** command line option is not accepted anymore.

B.3. Changes in binutils

Red Hat Developer Toolset 2.0 is distributed with **binutils 2.23.52**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version and the version

included in Red Hat Developer Toolset 1.1. Below is a comprehensive list of new features in this release.

The GNU assembler (**as**), GNU linker (**ld**), and other binary tools that are part of binutils are now released under the GNU General Public License, version 3.

B.3.1. GNU Linker

Another ELF linker, **gold**, is now available in addition to **ld**, the existing GNU linker. **gold** is intended to be a drop-in replacement for **ld**, so **ld**'s documentation is intended to be the reference documentation. **gold** supports most of **ld**'s features, except notable ones such as MRI-compatible linker scripts, cross-reference reports (**--cref**), and various other minor options. It also provides significantly improved link time with very large C++ applications.

In Red Hat Developer Toolset 2.0, the **gold** linker is not enabled by default. Users can explicitly switch between **ld** and **gold** by using the **alternatives** mechanism.

B.3.1.1. New Features

Changes Since Red Hat Enterprise Linux 6.4

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 6.4:

- ▶ A new **INPUT_SECTION_FLAGS** keyword has been added to the linker script language. This keyword can be used to select input sections by section header flags.
- ▶ A new **SORT_BY_INIT_PRIORITY** keyword has been added to the linker script language. This keyword can be used to sort sections by numerical value of the GCC *init_priority* attribute encoded in the section name.
- ▶ A new **SORT_NONE** keyword has been added to the linker script language. This keyword can be used to disable section sorting.
- ▶ A new linker-provided symbol, **__ehdr_start**, has been added. When producing ELF output, this symbol points to the ELF file header (and nearby program headers) in the program's memory image.

Changes Since Red Hat Enterprise Linux 5.9

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 5.9:

- ▶ GNU/Linux targets now support the **STB_GNU_UNIQUE** symbol binding, a GNU extension to the standard set of ELF symbol bindings. The binding is passed on to the dynamic linker, which ensures that in the entire process there is only one symbol with the given name and type in use.



Note

The implementation of this feature depends on capabilities only found in newer versions of the **glibc** library. Consequently, this feature is currently available in Red Hat Developer Toolset for Red Hat Enterprise Linux 6.

- ▶ A new command line option **--no-export-dynamic** has been added. This option can be used to undo the effect of the **-E** and **--export-dynamic** options.
- ▶ A new command line option **--warn-alternate-em** has been added. This option can be used to display a warning if an ELF format object file uses an alternate machine code.
- ▶ A new linker script function **REGION_ALIAS** has been added. This function can be used to create

alias names of memory regions.

- ▶ A new command line option **-Ttext-segment address** has been added for ELF-based targets. This option can be used to set the address of the first byte of the text segment.
- ▶ A new linker script command **INSERT** has been added. This command can be used to augment the default script.
- ▶ In a linker script input section, it is now possible to specify a file within an archive by using the **archive:file** syntax.
- ▶ The **--sort-common** command line option now accepts **ascending** and **descending** as optional arguments. This can be used to specify which sorting order to use.
- ▶ A new command line option **--build-id** has been added for ELF-based targets. This option can be used to generate a unique per-binary identifier embedded in a note section.
- ▶ A new command line option **--default-script=file_name** (or **-dT file_name**) has been added. This option can be used to specify a replacement for the built-in linker script.
- ▶ A new command line option **-Bsymbolic-functions** has been added. When creating a shared library, this option will cause references to global function symbols to be bound to the definitions with the shared library, if such exist.
- ▶ The new command line options **--dynamic-list-cpp-new** and **--dynamic-list-data** have been added, which can be used to modify the dynamic list.

B.3.1.2. Compatibility Changes

Changes Since Red Hat Enterprise Linux 6.4

The following compatibility changes have been made since the release of binutils included in Red Hat Enterprise Linux 6.4:

- ▶ The **--copy-dt-needed-entries** command line option is no longer enabled by default. Instead, **-no-copy-dt-needed-entries** is now the default option.
- ▶ Evaluation of linker script expressions has been significantly improved. Note that this can negatively affect scripts that rely on undocumented behavior of the old expression evaluation.

Changes Since Red Hat Enterprise Linux 5.9

The following compatibility changes have been made since the release of binutils included in Red Hat Enterprise Linux 5.9:

- ▶ The **--add-needed** command line option has been renamed to **--copy-dt-needed-entries** in order to avoid confusion with the **--as-needed** option.
- ▶ For GNU/Linux systems, the linker no longer processes any relocations made against symbols of the **STT_GNU_IFUNC** type. Instead, it emits them into the resulting binary for processing by the loader.



Note

The implementation of this feature depends on capabilities only found in newer versions of the **glibc** library. Consequently, this feature is currently available in Red Hat Developer Toolset for Red Hat Enterprise Linux 6.

- ▶ The **--as-needed** command line option has been adapted to link in a dynamic library in the following two cases:
 1. if the dynamic library satisfies undefined symbols in regular objects, and
 2. if the dynamic library satisfies undefined symbols in other dynamic libraries unless the library

is already found in a **DT_NEEDED** entry of one of the libraries that are already linked.

- The **-l:file_name** command line option now searches the library path for a file name called **file_name** without adding the **.a** or **.so** file extension.

B.3.2. GNU Assembler

B.3.2.1. New Features

Changes Since Red Hat Enterprise Linux 6.4

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 6.4:

- The GNU Assembler no longer requires double ampersands in macros.
- A new **--compress-debug-sections** command line option has been added to enable compression of DWARF debug information sections in the relocatable output file. Compressed debug sections are currently supported by the **readelf**, **objdump**, and **gold** tools, but not by **ld**.
- Support for **.bundle_align_mode**, **.bundle_lock**, and **.bundle_unlock** directives for x86 targets has been added..
- On x86 architectures, the GNU Assembler now allows **rep bsf**, **rep bsr**, and **rep ret** syntax.

Changes Since Red Hat Enterprise Linux 5.9

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 5.9:

- GNU/Linux targets now support **gnu_unique_object** as a value of the **.type** pseudo operation. This value can be used to mark a symbol as globally unique in the entire process.
- Support for the new discriminator column in the DWARF line table with a discriminator operand for the **.loc** directive has been added.
- The **.type** pseudo operation now accepts a type of **STT_GNU_IFUNC**. This can be used to indicate that if the symbol is the target of a relocation, its value should not be used. Instead, the function should be invoked and its result used as the value.
- A new pseudo operation **.cfi_val_encoded_addr** has been added. This pseudo operation can be used to record constant addresses in unwind tables without runtime relocation.
- A new command line option **-msse-check=[none|error|warning]** has been added for x86 targets.
- The **-a** command line option now accepts **g** as a valid sub-option. This combination can be used to enable assembly listings with additional information about the assembly, including the list of supplied command line options or the assembler version.
- A new command line option **-msse2avx** has been added for x86 targets. This option can be used to encode SSE instructions with VEX prefix.
- x86 targets now support the Intel XSAVE, EPT, MOVBE, AES, PCLMUL, and AVX/FMA instructions.
- New command line options **-march=cpu[, +extension...]**, **-mtune=cpu**, **-mmnemonic=[att|intel]**, **-msyntax=[att|intel]**, **-mindex-reg**, **-mnaked-reg**, and **-mold-gcc** have been added for x86 targets.
- New pseudo operations **.string16**, **.string32**, and **.string64** have been added. These pseudo operations be used to generate wide character strings.
- The i386 port now supports the SSE5 instruction set.
- A new pseudo operation **.reloc** has been added. This pseudo operation serves as a low-level

interface for creating relocations.

B.3.3. Other Binary Tools

B.3.3.1. New Features

Changes Since Red Hat Developer Toolset 1.1

The following features have been added since the release of binutils included in Red Hat Developer Toolset 1.1:

- ▶ A manual page for the **dwp** utility has been added.
- ▶ The binary tools now provide support for the AMD Family 15h processors, models 02h and 10-1fh.

Changes Since Red Hat Enterprise Linux 6.4

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 6.4:

- ▶ The **readelf** and **objdump** tools can now display the contents of the **.debug.macro** sections.
- ▶ New **--dwarf-start** and **--dwarf-end** command line options have been added to the **readelf** and **objdump** tools. These options are used by the new Emacs mode (see the **dwarf-mode.el** file).
- ▶ A new **--interleave-width** command line option has been added to the **objcopy** tool to allow the use of the **--interleave** to copy a range of bytes from the input to the output.
- ▶ A new **--dyn-syms** command line option has been added to the **readelf** tool. This option can be used to dump dynamic symbol table.
- ▶ A new tool, **elfedit**, has been added to binutils. This tool can be used to directly manipulate ELF format binaries.
- ▶ A new command line option **--addresses** (or **-a** for short) has been added to the **addr2line** tool. This option can be used to display addresses before function and source file names.
- ▶ A new command line option **--pretty-print** (or **-p** for short) has been added to the **addr2line** tool. This option can be used to produce human-readable output.
- ▶ Support for **dwz -m** optimized debug information has been added.
- ▶ The *devtoolset-2-binutils-devel* package now provides the **demangle.h** header file.

Changes Since Red Hat Enterprise Linux 5.9

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 5.9:

- ▶ A new command line option **--insn-width=width** has been added to the **objdump** tool. This option can be used to specify the number of bytes to be displayed on a single line when disassembling instructions.
- ▶ A new command line option **--relocated-dump=name|number** has been added to the **readelf** tool. This option can be used to display the relocated contents of a section as a sequence of bytes.
- ▶ A new command line option **--external-symbols-table=filename** has been added to the **gprof** tool. This option can be used to read a symbol table from a certain file.
- ▶ **bfd** now supports a plugin target, which can be used to get basic support for new file formats by having the plugin target load the same shared objects used by **gold**.
- ▶ The **--dwarf** (or **-W** for short) command line option of the **objdump** tool has been adapted to be as flexible as the **--debug-dump** (or **-w**) option of **readelf**.

- New command line options **--prefix=prefix** and **--prefix-strip=level** have been added to the **objdump** tool. These options can be used to add absolute paths for the **--source** (or **-S** for short) option.
- A new command line option **-wL** has been added to the **readelf** tool. This option can be used to dump decoded contents of the **.debug_line** section.
- “Thin” archives are now supported. Instead of containing object files, such archives contain just pathnames pointing to those files.
- A new command line option **-F** has been added to the **objdump** tool. This option can be used to include file offsets in the disassembly.
- A new command line option **-c** has been added to the **readelf** tool. This option can be used to allow string dumps of archive symbol index.
- The i386 port now supports the SSE5 instruction set.
- A new command line option **-p** has been added to the **readelf** tool. This option can be used to allow string dumps of sections.

B.3.3.2. Compatibility Changes

Changes Since Red Hat Enterprise Linux 5.9

The following compatibility changes have been made since the release of binutils included in Red Hat Enterprise Linux 5.9:

- The **--as-needed** command line option has been adapted to link in a dynamic library in the following two cases:
 1. if the dynamic library satisfies undefined symbols in regular objects, and
 2. if the dynamic library satisfies undefined symbols in other dynamic libraries unless the library is already found in a **DT_NEEDED** entry of one of the libraries that are already linked.

B.4. Changes in elfutils

Red Hat Developer Toolset 2.0 is distributed with **elfutils 0.155**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version and the version included in Red Hat Developer Toolset 1.1. Below is a comprehensive list of new features in this release.

B.4.1. Changes Since Red Hat Developer Toolset 1.1

The following features have been added since the release of elfutils included in Red Hat Developer Toolset 1.1:

- In the **libdw** library, the **DW_LANG_ObjC** constant has been correctly renamed to **DW_LANG_ObjC**. Note that any existing source code that uses the old name needs to be updated accordingly.
- The **libdw** library now supports new constants **DW_ATE_UTF** and **DW_OP_GNU_parameter_ref**. In addition, it also defines a family of constants to support the **.debug_macro** section, namely **DW_MACRO_GNU_define**, **DW_MACRO_GNU_undef**, **DW_MACRO_GNU_start_file**, **DW_MACRO_GNU_end_file**, **DW_MACRO_GNU_define_indirect**, **DW_MACRO_GNU_undef_indirect**, and **DW_MACRO_GNU_transparent_include**.
- When working with the **libelf** library, both **elf32_xlatetomd** and **elf64_xlatetomd** now work for cross-endian ELF note data.
- The **elf_getshdr()** function provided by the **libelf** library has been corrected to work consistently on non-mmapped ELF files after making the **elf_cntl(ELF_C_FDREAD)** function call.
- The **libelf** library now supports **ar** archives with a 64-bit symbol table.

- ▶ The **eu-readelf** command is now able to display the contents of the **.debug_macro** section.
- ▶ The **eu-readelf** command now correctly recognizes the **DW_OP_GNU_parameter_ref** DWARF GNU extension opcode in location expressions.

B.4.2. Changes Since Red Hat Enterprise Linux 6.4

The following features have been added since the release of elfutils included in Red Hat Enterprise Linux 6.4:

- ▶ The **libdw** library now handles compressed debuginfo sections. The **dwarf_highpc()** function now handles the DWARF 4 **DW_AT_high_pc** constant form.
- ▶ The **eu-elflint** utility now accepts executables produced by the **gold** linker.
- ▶ The **eu-nm** utility now supports C++ demangling.
- ▶ The **eu-ar** utility now supports a new modifier **D** for *deterministic output* with no UID, GID, or mtime info. The **U** modifier is the inverse.
- ▶ The **eu-readelf** utility can now print SDT ELF notes (*SystemTap probes*) and the **.gdb_index** GDB section. It can now also print **DW_OP_GNU_entry_value** and **DW_AT_GNU_call_site** families of DIE attributes.
- ▶ The **eu-strip** utility now recognizes a new command line option, **--reloc-debug-sections**.

B.4.3. Changes Since Red Hat Enterprise Linux 5.9

In addition to the above changes, the following features have been added since the release of elfutils included in Red Hat Enterprise Linux 5.9:

- ▶ DWARF 4 is now supported. As well, support for **DW_OP_GNU_implicit_pointer** and **STB_GNU_UNIQUE** has been added.
- ▶ The **libdwfl** library now supports automatic decompression of files compressed with **gzip**, **bzip2** and **lzma**, and of Linux kernel images made with **gzip**, **bzip2** or **lzma**. Files named with compression suffixes are searched for Linux kernel images. Core file support was improved. Support has been added for decoding DWARF CFI into location description form. Support has been added for some new DWARF 3 expression operations, which were previously omitted. A new function, **dwfl_dwarf_line()**, has been added.
- ▶ The **eu-elfcmp** utility now supports a new command line option **--ignore-build-id** to ignore differing build ID bits. The new option **--verbose** (or **-l** for short) prints all differences.
- ▶ The **eu-strip** utility now recognizes a new command line option **--strip-sections** to remove section headers entirely.
- ▶ The **libdw** library now has new functions **dwarf_next_unit()**, **dwarf_offdie_types()**, **dwarf_lineisa()**, **dwarf_linediscriminator()**, **dwarf_lineop_index()**, **dwarf_getlocation_implicit_pointer()**, and **dwarf_aggregate_size()**.
- ▶ The **eu-addr2line** utility now recognizes a new command line option **--flags** (or **-F** for short) to print more DWARF line information details.
- ▶ The **libelf** library now supports using more than 65536 program headers in a file. In addition, a new function **elf_getphdrnum()** has been added.
- ▶ The **eu-addr2line** utility now accepts the **--section=name** option (or **-j name** for short) in the interests of binutils compatibility.
- ▶ **libcpu** Intel SSE4 disassembler support has been added.
- ▶ The **eu-readelf** utility now implements call frame information and exception handling dumping. The **-e** command line option has been added; this is enabled implicitly by **-a**.

B.5. Changes in dwz

Red Hat Developer Toolset 2.0 is distributed with **dwz 0.11**, which provides a number of bug fixes and enhancements over the version included in Red Hat Developer Toolset 1.1. Below is a comprehensive list of changes in this release.

B.5.1. Changes Since Red Hat Developer Toolset 1.1

The following features have been added since the release of **dwz** included in Red Hat Developer Toolset 1.1:

- ▶ The **DW_FORM_data4**, **DW_FORM_data8**, and **DW_AT_high_pc** attributes have been optimized.
- ▶ The **dwz** utility now allocates more memory when recomputing abbreviations and no longer terminates unexpectedly.
- ▶ The **dwz** utility no longer crashes when processing debugging information that contains several copies of the same Debugging Information Entry (DIE) within the same compilation unit.
- ▶ The **dwz** utility no longer writes an incorrect **DW_FORM** code for a compilation unit version.
- ▶ The **dwz** utility now supports version 8 of the **.gdb_index** section.

B.6. Changes in GDB

Red Hat Developer Toolset 2.0 is distributed with **GDB 7.6**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version and the version included in Red Hat Developer Toolset 1.1. Below is a comprehensive list of new features in this release.

B.6.1. Changes Since Red Hat Developer Toolset 1.1

The following features have been added since the release of GDB included in Red Hat Developer Toolset 1.1:

- ▶ Target **record** has been renamed to **record-full**. Consequently, you can now use the **record full** command to record or replay an execution log. In addition, the following commands have been renamed:
 - The **set record insn-number-max** and **show record insn-number-max** commands have been renamed to **set record full insn-number-max** and **show record full insn-number-max**.
 - The **set record memory-query** and **show record memory-query** commands have been renamed to **set record full memory-query** and **show record full memory-query**.
 - The **set record stop-at-limit** and **show record stop-at-limit** commands have been renamed to **set record full stop-at-limit** and **show record full stop-at-limit**.
- ▶ A new record target, **record-btrace**, has been added. This target uses hardware support to record the control flow of a process and can be enabled by using the **record btrace** command. This record target does not support replaying the execution.



Important

The **record-btrace** target is only available on Intel Atom processors and requires the Linux kernel in version 2.6.32 or later.

- » New **record instruction-history** and **record function-call-history** commands have been added. These commands allow you to view information about an execution log without having to replay it. The **record instruction-history** command displays the execution history at instruction granularity and the **record function-call-history** displays the execution history at function granularity. The commands are only supported by the **record btrace** command.
- » A new command line option, **-nh**, has been added. This option allows you to disable automatic loading of the `~/.gdbinit` file without disabling other initialization files.
- » The **-epoch** command line option has been removed. This option was used by GDB mode in Epoch, a deprecated clone of the Emacs text editor.
- » The **ptype** and **whatis** commands have been updated to accept an argument to control the type formatting.
- » The **info proc** command has been updated to work on some core files.
- » The **cd** command has been enhanced and no longer requires a directory path as its first argument. When executed with no arguments, the command now changes to the home directory.
- » GDB now uses *GNU v3 ABI* as the default C++ ABI. This has been the default option for GCC since November 2000.
- » The **info tracepoints** command has been enhanced to display **installed on target** or **not installed on target** for each non-pending location of a tracepoint.
- » A new command, **fo**, has been added. This command serves as a shorter variant of the **forward-search** command.
- » A new command, **catch signal**, has been added. This command can be used to catch signals by their names and is similar to the **handle** command, but also allows you to attach additional conditions or commands.
- » A new command, **maint info bfds**, has been added. This command can be used to list all binary files (BFDs) opened by GDB.
- » Two new commands, **python-interactive [command]** and its shorter variant **pi [command]**, have been added. These commands allow you to start an interactive Python prompt or evaluate a Python command and print the results to standard output.
- » A new command, **py [command]**, has been added. This command serves as a shorter variant of the **python [command]** command.
- » New **enable type-printer [name...]** and **disable type-printer [name...]** commands have been added. These commands allow you to enable or disable type printers.
- » New **set print type methods on|off** and **show print type methods** commands have been added. These commands allow you to control whether method declarations are displayed by the **ptype** command. This functionality is enabled by default.
- » New **set print type typedefs on|off** and **show print type typedefs** commands have been added. These commands allow you to control whether **typedef** definitions are displayed by the **ptype** command. This functionality is enabled by default.
- » New **set filename-display basename|relative|absolute** and **show filename-display** commands have been added. These commands allow you to control the way in which file names are displayed: the **basename** option displays only the base name of a file name, **relative** displays a path relative to the compilation directory, and **absolute** displays an absolute path to the file. The default option is **relative** to preserve the previous behavior.
- » New **set trace-buffer-size** and **show trace-buffer-size** commands have been added. These commands allow you to control the size of the trace buffer for a target.
- » New **set remote trace-buffer-size-packet auto|on|off** and **show remote trace-buffer-size-packet** commands have been added. These commands allow you to control the

use of the remote protocol **QTBuffer:size** packet.

- ▶ New **set debug notification** and **show debug notification** commands have been added. These commands allow you to control whether to display debugging information for asynchronous remote notification. This functionality is disabled by default.
- ▶ New convenience functions **\$_memeq(buf1, buf2, length)**, **\$_streq(str1, str2)**, **\$_strlen(str)**, and **\$_regex(str, regex)** have been added.

The following changes have been made to the Python scripting support since the release of GDB included in Red Hat Developer Toolset 1.1:

- ▶ Users can now create vectors by using the **`gdb.Type.vector()`** method.
- ▶ The **`atexit.register()`** method is now supported.
- ▶ Users can now pretty-print types by using the Python API.
- ▶ In addition to **Python 2.4** and later, GDB now also supports **Python 3**.
- ▶ A new class, **`gdb.Architecture`**, has been added. This class exposes the internal representation of the architecture in the Python API.
- ▶ A new method, **`Frame.architecture`**, has been added. This method can be used to return the **`gdb.Architecture`** object corresponding to the frame's architecture.
- ▶ Frame filters and frame decorators have been added.

The following MI changes have been made since the release of GDB included in Red Hat Developer Toolset 1.1:

- ▶ A new async record, **`=cmd-param-changed`**, has been added. This async record reports that a command parameter has changed.
- ▶ A new async record, **`=traceframe-changed`**, has been added. This async record reports that a trace frame has been changed by using the **`tfind`** command.
- ▶ New async records **`=tsv-created`**, **`=tsv-deleted`**, and **`=tsv-modified`** have been added. These async records report that a trace state variable has been created, deleted, or modified.
- ▶ New async records **`=record-started`** and **`=record-stopped`** have been added. These async records report that a process record has been started or stopped.
- ▶ A new async record, **`=memory-changed`**, has been added. This async record reports that the memory has changed.
- ▶ When the source is requested, the **`-data-disassemble`** command now includes a new **`fullname`** field containing an absolute path to the source file name.
- ▶ A new optional parameter, **`COUNT`**, has been added to the **`-data-write-memory-bytes`** command. This parameter can be used to allow pattern filling of memory areas.
- ▶ New commands **`-catch-load`** and **`-catch-unload`** have been added. These commands can be used to intercept shared library **`load/unload`** events.
- ▶ The response to breakpoint commands and breakpoint async records now includes a new **`installed`** field. This field reports the current state of each non-pending tracepoint location: when the tracepoint is installed, the value of this field is **`y`**, otherwise the value is **`n`**.
- ▶ The output of the **`-trace-status`** command now includes a new **`trace-file`** field. This field is only present when examining a trace file and contains the name of this file.
- ▶ The **`fullname`** field is now always present along with the **`file`** field. This field is included even if GDB cannot find the file.

A number of new remote packets have been added since the release of GDB included in Red Hat Developer Toolset 1.1. See [Table B.3, “New Remote Packets”](#) for a complete list.

Table B.3. New Remote Packets

| Remote Packet | Description |
|---|---|
| QTBUFFER:SIZE | Sets the size of the trace buffer. The remote stub reports support for this packet to the qSupported query. |
| Qbtrace:bts | Enables branch tracing based on <i>Branch Trace Store</i> (BTS) for the current thread. The remote stub reports support for this packet to the qSupported query. |
| Qbtrace:off | Disables branch tracing for the current thread. The remote stub reports support for this packet to the qSupported query. |
| qXfer:btrace:read | Reads the traced branches for the current thread. The remote stub reports support for this packet to the qSupported query. |
| qXfer:libraries-svr4:read 's annex | <p>The previously unused annex of the qXfer:libraries-svr4:read packet is now used to support passing of an argument list. The remote stub reports support for this argument list to the qSupported query.</p> <p>The defined arguments are start and prev. These arguments are used to reduce work necessary for updating the library list and significantly speed up the process.</p> |

B.6.2. Changes Since Red Hat Enterprise Linux 6.4

The features below have been added since the release of GDB included in Red Hat Enterprise Linux 6.4.

New Features

- ▶ Support for linespecs has been improved (in particular, a more consistent handling of ambiguous linespecs, some support for labels in the program's source, and FILE:LINE support now extends to further linespecs types). Breakpoints are now set on all matching locations in all inferiors and will be updated according to changes in the inferior.
- ▶ New inferior control commands **skip function** and **skip file** have been added. These commands can be used to skip certain functions and files when stepping.
- ▶ The **info threads** command now displays the thread name as set by **prctl** or **pthread_setname_np**. In addition, new commands **thread name** and **thread find** have been added. The **thread name** command accepts a name as an argument and can be used to set the name of the current thread. The **thread find** command accepts a regular expression and allows the user to find threads that match it.
- ▶ GDB now provides support for reading and writing a new **.gdb_index** section. The command **gdb-add-index** can be used to add **.gdb_index** to a file, which allows GDB to load symbols from that file faster. Note that this feature is already present in Red Hat Enterprise Linux 6.1 and later.
- ▶ The **watch** command has been adapted to accept **-location** as an optional argument.
- ▶ Two new special values can now be used when specifying the current search path for **libthread_db**: **\$sdir** represents the default system locations of shared libraries, and **\$pdir** stands for the directory with the **libthread** that is used by the application.
- ▶ A new command **info macros** has been added. This command accepts linespec as an optional argument and can be used to display the definitions of macros at that linespec location. Note that in order to do this, the debugged program must be compiled with the **-g3** command line option to have

macro information available in it.

- ▶ A new command **alias** has been added. This command can be used to create an alias of an existing command.
- ▶ The **info macro** command now accepts **-all** and **--** as valid options.
- ▶ To display a function parameter's entry value (that is, the value at the time of function entry), the suffix **@entry** can be added to the parameter. GDB now displays **@entry** values in backtraces, if available.
- ▶ New **set print entry-values** and **show print entry-values** commands have been added. The **set print entry-values** command accepts **both**, **compact**, **default**, **if-needed**, **no**, **only**, and **preferred** as valid arguments and can be used to enable printing of function arguments at function entry. The **show print entry-values** command can be used to determine whether this feature is enabled.
- ▶ New **set debug entry-values** and **show debug entry-values** commands have been added. The **set debug entry-values** command can be used to enable printing of debugging information for determining frame argument values at function entry and virtual tail call frames.
- ▶ **!command** has been added as an alias of **shell command**.
- ▶ The **watch** command now accepts **mask mask_value** as an argument. This can be used to create masked watchpoints.
- ▶ New **set extended-prompt** and **show extended-prompt** commands have been added. The **set extended-prompt** command enables support for a defined set of escape sequences that can be used to display various information. The **show extended-prompt** command can be used to determine whether the extended prompt is enabled.
- ▶ New **set basenames-may-differ** and **show basenames-may-differ** commands have been added. The **set basenames-may-differ** command enables support for source files with multiple base names. The **show basenames-may-differ** command can be used to determine whether this support is enabled. The default option is **off** to allow faster GDB operations.
- ▶ A new command line option **-ix** (or **--init-command**) has been added. This option acts like **-x** (or **--command**), but is executed before loading the debugged program.
- ▶ A new command line option **-iex** (or **--init-eval-command**) has been added. This option acts like **-ex** (or **--eval-command**), but is executed before loading the debugged program.
- ▶ The **info os** command has been changed and can now display information on several objects managed by the operating system, in particular:
 - The **info os procggroups** command lists process groups.
 - The **info os files** command lists file descriptors.
 - The **info os sockets** command lists internet-domain sockets.
 - The **info os shm** command lists shared-memory regions.
 - The **info os semaphores** command lists semaphores.
 - The **info os msg** command lists message queues.
 - The **info os modules** command lists loaded kernel modules.
- ▶ GDB now has support for *Static Defined Tracing* (SDT) probes. Currently, the only implemented back end is for SystemTap probes (the **sys/sdt.h** header file). You can set a breakpoint by using the new **-probe**, **-pstatp**, or **-probe-statp** options, and inspect the probe arguments by using the new **\$_probe_arg** family of convenience variables.
- ▶ The **symbol-reloading** option has been deleted.
- ▶ **gdbserver** now supports STDIO connections, for example:

```
(gdb) target remote | ssh myhost gdbserver - hello
```

- GDB is now able to print *flag* enums. In a flag enum, all enumerator values have no bits in common when pairwise AND-ed. When GDB prints a value whose type is a flag enum, GDB shows all the constants; for example, for enum **E { ONE = 1, TWO = 2 }**:

```
(gdb) print (enum E) 3
$1 = (ONE | TWO)
```

- The file name part of a linespec now matches trailing components of a source file name. For example, **break gcc/expr.c:1000** now sets a breakpoint in the **build/gcc/expr.c** file, but not in **build/libcpp/expr.c**.
 - The **info proc** and **generate-core-file** commands now work on remote targets connected to **gdbserver**.
 - The command **info catch** has been removed.
 - The Ada-specific **catch exception** and **catch assert** commands now accept conditions at the end of the command.
 - The **info static-tracepoint-marker** command now works on native targets with an in-process agent.
 - GDB can now set breakpoints on inline functions.
 - The **.gdb_index** section has been updated to include symbols for inline functions. By default, GDB now ignores older **.gdb_index** sections until their **.gdb_index** sections can be recreated. The new command **set use-deprecated-index-sections on** causes GDB to use any older **.gdb_index** sections it finds. If this option is set, the ability to set breakpoints on inline functions is lost in symbol files with older **.gdb_index** sections.
- The **.gdb_index** section has also been updated to record more information about each symbol.
- GDB now provides Ada support for GDB/MI Variable Objects.
 - GDB now supports **breakpoint always-inserted mode** in the **record** target.
 - **gdbserver** now supports evaluation of breakpoint conditions. Note that you can instruct GDB to send the breakpoint conditions in bytecode form, but **gdbserver** only reports the breakpoint trigger to GDB when its condition evaluates to true.
 - The **z0/z1** breakpoint insertion packets have been extended to carry a list of conditional expressions over to the remote stub depending on the condition evaluation mode. You can use the **set remote conditional-breakpoints-packet** command to control the use of this extension.
 - A new RSP packet **QProgramSignals** can be used to specify the signals the remote stub can pass to the debugged program without GDB involvement.
 - A new command **-info-os** has been added as the MI equivalent of **info os**.
 - Output logs, such as **set logging** and related, now include MI output.
 - New **set use-deprecated-index-sections on|off** and **show use-deprecated-index-sections on|off** commands have been added. These commands allow you to control the use of deprecated **.gdb_index** sections.
 - New **catch load** and **catch unload** commands have been added. These commands allow you to stop execution of a debugged program when a shared library is loaded or unloaded.
 - A new command **enable count** has been added. This command allows you to auto-disable a breakpoint after several hits.
 - A new command **info vtbl** has been added. This command allows you to show the virtual method tables for C++ and Java objects.

- ▶ A new command **explore** has been added. It supports two subcommands **explore value** and **explore type**, and allows you to recursively explore values and types of expressions. Note that this command is only available with Python-enabled GDB.
- ▶ A new command **dprintf location,format,args...** has been added. This command allows you to create a dynamic **printf**-type breakpoint, which performs a **printf**-like operation and then resumes program execution.
- ▶ New **set print symbol** and **show print symbol** commands have been added. These commands allow you to control whether GDB attempts to display the symbol, if any, that corresponds to addresses it prints. This functionality is enabled by default, but you can restore the previous behavior by running the **set print symbol off** command.
- ▶ New **set breakpoint condition-evaluation** and **show breakpoint condition-evaluation** commands have been added. These commands allow you to control whether breakpoint conditions are evaluated by GDB (the **host** option), or by **gdbserver** (the **target** option). The default option, **auto**, chooses the most efficient available mode.
- ▶ New **set dprintf-style gdb|call|agent** and **show dprintf-style** commands have been added. These commands allow you to control the way in which a dynamic **printf** is performed: the **gdb** option requests a GDB **printf** command, **call** causes **dprintf** to call a function in the inferior, and **agent** requests that the target agent such as **gdbserver** does the printing.
- ▶ New **set dprintf-function expression**, **show dprintf-function**, **set dprintf-channel expression**, and **show dprintf-channel** commands have been added. These commands allow you to set the function and optional first argument to the call when using the **call** style of dynamic **printf**.
- ▶ New **set disconnected-dprintf on|off** and **show disconnected-dprintf** commands have been added. These commands allow you to control whether agent-style dynamic **printfs** continue to be in effect after GDB disconnects.

The following changes have been made to the C++ language support since the release of the GNU Debugger included in Red Hat Enterprise Linux 6.2:

- ▶ When debugging a template instantiation, parameters of the template are now put in scope.

The following changes have been made to the Python scripting support since the release of the GNU Debugger included in Red Hat Enterprise Linux 6.2:

- ▶ The **register_pretty_printer** function in module **gdb.printing** now takes an optional **replace** argument.
- ▶ The **maint set python print-stack on|off** command has been deprecated and will be deleted in GDB 7.5. The new command **set python print-stack none|full|message** has replaced it.
- ▶ A prompt substitution hook (**prompt_hook**) is now available to the Python API.
- ▶ A new Python module **gdb.prompt** has been added to the GDB Python modules library.
- ▶ Python commands and convenience-functions located in **data_directory/python/gdb/command/** and **data_directory/python/gdb/function/** are now automatically loaded on GDB start-up.
- ▶ Blocks now provide four new attributes: **global_block**, **static_block**, **is_static**, and **is_global**.
- ▶ The **gdb.breakpoint** function has been deprecated in favor of **gdb.breakpoints**.
- ▶ A new class **gdb.FinishBreakpoint** is provided.
- ▶ Type objects for **struct** and **union** types now allow access to the fields using standard Python

dictionary (mapping) methods.

- ▶ A new event **`gdb.new_objfile`** has been added.
- ▶ A new function **`deep_items`** has been added to the **`gdb.types`** module.
- ▶ The function **`gdb.Write`** now accepts an optional keyword **`stream`**.
- ▶ Parameters can now be sub-classed in Python, which allows for implementation of the **`get_set_doc`** and **`get_show_doc`** functions.
- ▶ Symbols, Symbol Table, Symbol Table and Line, Object Files, Inferior, Inferior Thread, Blocks, and Block Iterator APIs now have an **`is_valid`** method.
- ▶ Breakpoints can now be sub-classed in Python, which allows for implementation of the **`stop`** function that is executed each time the inferior reaches that breakpoint.
- ▶ A new function **`gdb.lookup_global_symbol`** has been added. This function can be used to look up a global symbol.
- ▶ GDB values in Python are now callable if the value represents a function.
- ▶ A new module **`gdb.types`** has been added.
- ▶ A new module **`gdb.printing`** has been added.
- ▶ New commands **`info pretty-printers`**, **`enable pretty-printer`**, and **`disable pretty-printer`** have been added.
- ▶ A new **`gdb.parameter("directories")`** function call is now available.
- ▶ A new function **`gdb.newest_frame`** has been added. This function can be used to return the newest frame in the selected thread.
- ▶ The **`gdb.InferiorThread`** class now supports a new **`name`** attribute.
- ▶ Support for inferior events has been added. Python scripts can now add observers in order to be notified of events occurring in the process being debugged.
- ▶ GDB commands implemented in Python can now be put in the **`gdb.COMMAND_USER`** command class.
- ▶ The **`maint set python print-stack on|off`** command has been removed and replaced by **`set python print-stack`**.
- ▶ A new class **`gdb.printing.FlagEnumerationPrinter`** has been added. This class can be used to apply **`flag enum`**-style pretty-printing to enums.
- ▶ The **`gdb.lookup_symbol`** function now works correctly when there is no current frame.
- ▶ The **`gdb.Symbol`** object now has an additional attribute **`line`**. This attribute holds the line number in the source at which the symbol was defined.
- ▶ The **`gdb.Symbol`** object now has an additional attribute **`needs_frame`**, and a new method **`value`**. The **`needs_frame`** attribute indicates whether the symbol requires a frame to compute its value, and the **`value`** method computes the symbol's value.
- ▶ The **`gdb.Value`** object now has a new method **`referenced_value`**. This method can be used to dereference a pointer as well as C++ reference values.
- ▶ The **`gdb.Symtab`** object now has two new methods, **`global_block`** and **`static_block`**. These methods return the global and static blocks (as **`gdb.Block`** objects) of the underlying symbol table respectively.
- ▶ A new method **`gdb.find_pc_line`** returns the **`gdb.Symtab_and_line`** object associated with a PC value.
- ▶ The **`gdb.Symtab_and_line`** object now has an additional attribute **`last`**. This attribute holds the end of the address range occupied by the code for the current source line.

Compatibility Changes

- A new command **info auto-load** has been added and can be used to display the status of various automatically loaded files. The **info auto-load gdb-scripts** command lists automatically loaded canned sequences of commands, **info auto-load python-scripts** displays the status of automatically loaded Python scripts, **info auto-load local-gdbinit** displays whether a local **.gdbinit** file in the current working directory is loaded, and **info auto-load libthread-db** displays whether the inferior-specific thread debugging shared library is loaded.
- New commands **set auto-load** and **show auto-load** have been added and can be used to control automatic loading of files:
 - The **set auto-load gdb-scripts** and **show auto-load gdb-scripts** commands control automatic loading of GDB scripts.
 - The **set auto-load python-scripts** and **show auto-load python-scripts** commands control automatic loading of Python scripts.
 - The **set auto-load local-gdbinit** and **show auto-load local-gdbinit** commands control automatic loading of **.gdbinit** from the current working directory.
 - The **set auto-load libthread-db** and **show auto-load libthread-db** commands control automatic loading of inferior-specific **libthread_db**.
 - The **set auto-load scripts-directory** and **show auto-load scripts-directory** commands control the list of directories from which to automatically load GDB and Python scripts.
 - The **set auto-load safe-path** and **show auto-load safe-path** commands control the list of directories from which it is safe to automatically load all previously mentioned items.
 - The **set debug auto-load** and **show debug auto-load** commands control displaying of debugging information for all previously mentioned items.
- The **set auto-load off** command can be used to disable automatic loading globally. You can also use **show auto-load** with no subcommand to display current settings of all previously mentioned items.
- The **maint set python auto-load on|off** command has been replaced with **set auto-load python-scripts on|off**.
- The **maintenance print section-scripts** command has been renamed to **info auto-load python-scripts [pattern]** and is no longer classified as a maintenance-only command.
- Support for the Guile extension language has been removed.
- The GNU Debugger has been adapted to follow GCC's rules on accessing volatile objects when reading or writing target state during expression evaluation.

B.6.3. Changes Since Red Hat Enterprise Linux 5.9

In addition to the above changes, the features below have been added since the release of GDB included in Red Hat Enterprise Linux 5.9.

New Features

- For remote targets, debugging of shared libraries is now supported by default.
- New commands **set observer** and **show observer** have been added. The **set observer** command accepts **on** or **off** as an argument and can be used to allow or disallow the GNU Debugger to affect the execution of the debugged program. Use the **show observer** command to determine whether observer mode is enabled.
- A new convenience variable **\$_thread** has been added. This variable stores the number of the current thread.
- The **source** command now accepts **-s** as a valid option. This option can be used to search for the

script in the source search path regardless of the path in the file name.

- » Support for tracepoints, including fast and static tracepoints, has been added to **gdbserver**.
- » The **--batch** command line option has been adapted to disable pagination and queries.
- » Direct support for the reading and writing byte, word, and double-word x86 general purpose registers such as **\$al** has been added.
- » The **commands** command now accepts a range of breakpoints as an argument.
- » The **rbreak** command now accepts a file name as part of its argument. This can be used to limit the functions selected by the supplied regular expression to those that are defined in the specified file.
- » Support for multi-program (sometimes referred to as multi-executable or multi-exec) debugging has been added. In particular, the GNU Debugger now supports the following commands:
 - The **add-inferior** command can be used to add a new inferior.
 - The **clone-inferior** command can be used to create a copy of an inferior with the same executable loaded.
 - The **remove-inferior** command accepts an inferior ID as an argument and can be used to remove an inferior.
- » Support for trace state variables has been added. In particular, the GNU Debugger now supports the following commands:
 - The **tvariable \$variable_name [= expression]** command can be used to define or modify a trace state variable.
 - The **info tvariables** command can be used to display a list of currently defined trace state variables and their values.
 - The **delete tvariable \$variable_name...** command can be used to delete one or more trace state variables.
- » A new **ftrace** has been added. This command accepts a function name, a line number, or an address as an argument, and can be used to define a fast tracepoint at that location.
- » Support for disconnected tracing, trace files, and circular trace buffer has been added.
- » A new **teval** command has been added. This command accepts one or more expressions to evaluate at a tracepoint.
- » The GNU Debugger has been adapted to parse the **0b** prefix of binary numbers exactly the same way as the GNU Compiler Collection.
- » The GNU Debugger now supports the following commands for process record and replay:
 - New commands **set record memory** and **show record memory** have been added. The **set record memory** command accepts **on** or **off** as an argument and can be used to enable or disable stopping the inferior when a memory change of the next instruction cannot be recorded. Use the **show record memory-query** command to determine whether this feature is enabled.
 - A new command **record save** has been added. This command accepts a file name as an argument and can be used to save the execution log to a file.
 - A new command **record restore** has been added. This command accepts a file name as an argument and can be used to restore the execution log from a file.
- » A new command **eval** has been added. This command accepts a format string followed by one or more arguments, transforms it to a command, and then executes it.
- » A new command **save breakpoints** has been added. This command accepts a file name as an argument and can be used to store all currently defined breakpoints to a file. To restore the saved breakpoints from this file, use the **source** command.
- » New commands **set may-write-registers**, **set may-write-memory**, **set may-insert-breakpoints**, **set may-insert-tracepoints**, **set may-insert-fast-tracepoints**, and

set may-interrupt have been added. All of these commands accept either **on** or **off** as an argument, and can be used to set individual permissions for the target.

- ▶ A new command **main info program-spaces** has been added. This command can be used to display information about currently loaded program spaces.
- ▶ New commands **set remote interrupt-sequence** and **show remote interrupt-sequence** have been added. The **set remote interrupt-sequence** command accepts **Ctrl-C**, **BREAK**, and **BREAK-g** as valid arguments, and can be used to specify which interrupt sequence to send to the remote target in order to interrupt its execution. Use the **show remote interrupt-sequence** to determine the current setting.
- ▶ New commands **set remote interrupt-on-connect** and **show remote interrupt-on-connect** have been added. The **set remote interrupt-on-connect** accepts either **on** or **off** as an argument, and can be used to enable sending an interrupt sequence to the remote target when the GNU Debugger connects to it. Use the **show remote interrupt-on-connect** command to determine whether this feature is enabled.
- ▶ The **set remotebreak** and **show remotebreak** commands have been deprecated and users are advised to use **set remote interrupt-sequence** and **show remote interrupt-sequence** instead.
- ▶ The **disassemble** command has been adapted to accept two arguments in the form of **start,+length**.
- ▶ The **source** command can now be used to read commands from Python scripts.

The following changes have been made to the C++ language support since the release of the GNU Debugger included in Red Hat Enterprise Linux 5.8:

- ▶ Argument-dependent lookup (ADL) now directs function search to the namespaces of its arguments regardless of whether the namespace has been imported.
- ▶ In addition to member operators, the GNU Debugger can now look up operators that are:
 - defined in the global scope,
 - defined in a namespace and imported via the **using** directive,
 - implicitly imported from an anonymous namespace, or
 - the argument-dependent lookup (ADL) operators.
- ▶ Support for printing of static const class members that are initialized in the class definition has been enhanced.
- ▶ Support for importing of namespaces has been added.
- ▶ The C++ expression parser has been adapted to handle the cast operators **static_cast<>**, **dynamic_cast<>**, **const_cast<>**, and **reinterpret_cast<>**.

The following changes have been made to the Python scripting support since the release of the GNU Debugger in Red Hat Enterprise Linux 5.8:

- ▶ The GNU Debugger is now installed with a new directory located at **/opt/rh/devtoolset-2/root/usr/share/gdb/python/**. This directory serves as a standard location for Python scripts written for GDB.
- ▶ The Python API has been adapted to provide access to symbols, symbol tables, program spaces, breakpoints, inferiors, threads, and frame's code blocks. Users are now also allowed to create custom GDB parameters from the API and manipulate them by using the **set** and **show** commands.
- ▶ New functions **gdb.target_charset**, **gdb.target_wide_charset**, **gdb.progspaces**, **gdb.current_progspace**, and **gdb.string_to_argv** have been added.
- ▶ A new exception **gdb.GdbError** has been added.

- ▶ The GNU Debugger now searches pretty-printers in the current program space.
- ▶ The GNU Debugger can now enable or disable pretty-printers individually.
- ▶ The GNU Debugger has been adapted to look for names of Python scripts to automatically load in a special section named `.debug_gdb_scripts`.

B.7. Changes in strace

Red Hat Developer Toolset 2.0 is distributed with **strace 4.7**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version. Below is a comprehensive list of new features in this release.

B.7.1. Changes Since Red Hat Enterprise Linux 6.4 and 5.9

The following features have been added since the release of strace in Red Hat Enterprise Linux 6.4 and 5.9:

- ▶ A new command line option, **-y**, has been added. This option can be used to print file descriptor paths.
- ▶ A new command line option, **-P**, has been added. This option can be used to filter system calls based on the file descriptor paths.
- ▶ A new command line option, **-I**, has been added. This option can be used to control how interactive **strace** is.
- ▶ A new command line utility, **strace-log-merge**, has been added. This utility can be used to merge timestamped **strace** output into a single file.
- ▶ The **strace** utility now uses optimized interfaces to extract data from the traced process for better performance.
- ▶ The **strace** utility now provides improved support for decoding of arguments for various system calls. In addition, a number of new system calls are supported.

B.8. Changes in SystemTap

Red Hat Developer Toolset 2.0 is distributed with **SystemTap 2.1**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version and the version included in Red Hat Developer Toolset 1.1. Below is a comprehensive list of new features in this release.

B.8.1. Changes Since Red Hat Developer Toolset 1.1

The following features have been added since the release of SystemTap included in Red Hat Developer Toolset 1.1:

- ▶ SystemTap has been updated to provide experimental support for **Dyninst**-based probing in the user space. Users can now execute the **stap** utility with the **--dyninst** command line option to instrument their own programs without the need to acquire **root** privileges.
- ▶ SystemTap is now distributed with a number of manual pages that provide a detailed explanation of common errors:
 - The **error::buildid(7stap)** manual page documents build-id verification failures.
 - The **error::dwarf(7stap)** manual page documents common DWARF debuginfo quality problems.
 - The **error::fault(7stap)** manual page documents memory access faults.
 - The **error::inode-uprobes(7stap)** manual page documents current limitations of inode-uprobes.

- The **error::pass1**(7stap) manual page documents pass 1 (parsing) errors.
- The **error::pass2**(7stap) manual page documents pass 2 (elaboration) errors.
- The **error::pass3**(7stap) manual page documents pass 3 (translation) errors.
- The **error::pass4**(7stap) manual page documents pass 4 (compilation) errors.
- The **error::pass5**(7stap) manual page documents pass 5 (execution) errors.
- The **error::process-tracking**(7stap) manual page documents user-space process tracking errors.
- The **error::reporting**(7stap) manual page documents how to report SystemTap bugs.
- The scripting language used by SystemTap now supports the **=~** operator. Users can use this operator to match regular expressions.
- The preprocessor used by the **stap** utility now supports an experimental macro facility. Users can use the following construct to define new macros:

```
@define name(parameter...) %( body %)
```

Refer to the *PREPROCESSOR MACROS* section of the **stap**(1) manual page for more information on how to use this feature.

- The backtrace-related tapset functions have been improved and standardized.

B.9. Changes in OProfile

Red Hat Developer Toolset 2.0 is distributed with **OProfile 0.9.8**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version and the version included in Red Hat Developer Toolset 1.1. Below is a comprehensive list of new features in this release.

B.9.1. Changes Since Red Hat Developer Toolset 1.1

The following features have been added since the release of OProfile included in Red Hat Developer Toolset 1.1:

- A new command line utility, **operf**, has been added. This utility can be executed by a non-**root** user and can be used to collect profiling data.



Important

The **operf** utility relies on a kernel feature that is currently not available in Red Hat Enterprise Linux 5. Consequently, this utility only works and is supported on Red Hat Enterprise Linux 6.

- OProfile now supports Intel processors code-named Ivy Bridge.

B.9.2. Changes Since Red Hat Enterprise Linux 5.9

The following features have been added since the release of OProfile included in Red Hat Enterprise Linux 5.9:

- OProfile now supports AMD family11h, family12h, family14h, and AMD family15h processors.
- Generation of XML output has been corrected.
- Handling of the **--session-dir** command line option has been improved and a possible buffer overflow in the XML generator has been fixed.

B.10. Changes in Valgrind

Red Hat Developer Toolset 2.0 is distributed with **Valgrind 3.8.1**, which provides a number of bug fixes over the version included in Red Hat Developer Toolset 1.1. Below is a comprehensive list of changes in this release.

B.10.1. Changes Since Red Hat Developer Toolset 1.1

The following features have been added since the release of Valgrind included in Red Hat Developer Toolset 1.1:

- ▶ The DWARF debug information for Valgrind is now distributed in a separate package, *devtoolset-2-valgrind-debuginfo*.
- ▶ The *SGCHECK OPTIONS* section of the **valgrind**(1) manual page has been corrected and no longer includes the BBV options. As well, the BBV options are now listed in a separate section.

B.10.2. Changes Since Red Hat Enterprise Linux 5.9

In addition to the above changes, the following features have been added since the release of Valgrind 3.5.0 included in Red Hat Enterprise Linux 5.9:

- ▶ When running in 64-bit mode, Valgrind now supports the SSE4.2 instruction set with the exception of SSE4.2 AES instructions. In 32-bit mode, Valgrind only provides support up to and including the SSSE3 instruction set.
- ▶ A new processing script **cg_diff** has been added to **Cachegrind**. This processing script can find the difference between two profiles, and can therefore be used to evaluate the performance effects of a change in a program.
- ▶ The behavior of the **cg_annotate**'s rarely-used **--threshold** option has been changed.
- ▶ **Callgrind** now supports branch prediction simulation and can optionally count the number of executed global bus events. Note that in order to use this functionality for a better approximation of a "Cycle Estimation" as a derived event, you must manually update the event formula in **KCachegrind**.
- ▶ To accommodate machines with three levels of caches, both **Cachegrind** and **Callgrind** now refer to the LL (last-level) cache rather than the L2 cache. When **Cachegrind** or **Callgrind** auto-detects the cache configuration of such a machine, it now runs the simulation as if the L2 cache is not present. Consequently, the results are less likely to match the true result for the machine and should not be considered authoritative, but provide a general idea about a program's locality.
- ▶ A new command line option **--pages-as-heap** has been added to **Massif**. When this option is enabled, **Massif** tracks memory allocations at the level of memory pages (as mapped by **mmap**, **brk** and similar functions) instead of tracking allocations at the level of heap blocks (as allocated with **malloc**, **new**, or **new[]**). Each mapped page is treated as its own block. Interpreting the page-level output is harder than the heap-level output, but allows you to account for every byte of memory used by a program. By default, the **--pages-as-heap** option is disabled.
- ▶ New command line options **--free-is-write** and **--trace-alloc** have been added to **DRD**. The **--free-is-write** option allows you to detect reading from already freed memory, the **--trace-alloc** can be used to trace of all memory allocations and deallocations.
- ▶ A number of new allocations have been added to **DRD**. As well, you can now annotate custom barrier implementations and benign races on static variables.
- ▶ **DRD**'s happens before and happens after annotations have been enhanced and can be used to annotate, for example, a smart pointer implementation.
- ▶ **Helgrind**'s annotation set has been significantly improved to provide a general set of annotations to describe locks, semaphores, barriers, and condition variables. In addition, **Helgrind** now supports

annotations to describe thread-safe reference counted heap objects.

- » A new command line option **--show-possibly-lost** has been added to **Memcheck**. By default this option is enabled and causes the leak detector to show possibly-lost blocks.
- » A new experimental heap profiler, **DHAT** (Dynamic Heap Analysis Tool), has been added. **DHAT** keeps track of allocated heap blocks, and also inspects every memory reference to see which block (if any) is being accessed. This gives a lot of insight into block lifetimes, utilization, turnover, liveness, and the location of hot and cold fields. You can use **DHAT** to do hot-field profiling.
- » Support for unfriendly self-modifying code has been improved, and the extra overhead incurred by **--smc-check=all** has been reduced by approximately a factor of 5 as compared with the previous version of Valgrind.
- » A new command line option **--fullpath-after** has been added. This option can be used to display directory names for source files in error messages, and is combined with a flexible mechanism for specifying which parts of the paths should be shown.
- » A new command line option **--require-text-symbol** has been added. This option stops the execution if a specified symbol is not found in a given shared object when loaded into the process. As a result, working with function intercepting and wrapping is now safer and more reliable.
- » Valgrind now implements more reliable stack unwinding on amd64-linux, particularly in the presence of function wrappers, and with gcc-4.5 compiled code.
- » Valgrind now implements modest scalability (performance improvements) for very large, long-running applications.
- » Valgrind now provides improved support for analyzing programs that are running in **Wine**. Users can now include the **valgrind/valgrind.h**, **valgrind/memcheck.h**, and **valgrind/drd.h** header files in Windows programs that are compiled with MinGW or one of the Microsoft Visual Studio compilers.

Revision History

| | | |
|---|------------------------|-------------------------|
| Revision 2.0-2 | Tue 11 Mar 2014 | Jaromír Hradílek |
| Red Hat Developer Toolset 2.1 GA release of the <i>User Guide</i> . | | |
| Revision 1.0-3 | Tue 11 Feb 2014 | Jaromír Hradílek |
| Red Hat Developer Toolset 2.1 Beta release of the <i>User Guide</i> | | |
| Revision 1.0-2 | Tue 10 Sep 2013 | Jaromír Hradílek |
| Red Hat Developer Toolset 2.0 GA release of the <i>User Guide</i> . | | |
| Revision 1.0-1 | Tue 06 Aug 2013 | Jaromír Hradílek |
| Red Hat Developer Toolset 2.0 Beta-2 release of the <i>User Guide</i> . | | |
| Revision 1.0-0 | Tue 28 May 2013 | Jaromír Hradílek |
| Red Hat Developer Toolset 2.0 Beta-1 release of the <i>User Guide</i> . | | |

Index

A

ABI

- compatibility, [ABI Compatibility](#)

addr2line

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

application binary interface (see ABI)

ar

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

as (see GNU assembler)

assembling (see GNU assembler)

B

bfd

- features, [New Features](#)

binutils

- documentation, [Additional Resources](#)
- features, [Main Features](#)
- installation, [Installing binutils](#)
- overview, [binutils](#)

- usage, [Using the GNU Assembler](#), [Using the GNU Linker](#), [Using Other Binary Tools](#)
- version, [About Red Hat Developer Toolset](#), [binutils](#)

C

C programming language

- compiling, [Using the C Compiler](#), [Preparing a Program for Debugging](#)
- running, [Running a C Program](#)
- support, [GNU C Compiler](#)

C++ programming language

- compiling, [Using the C++ Compiler](#), [Preparing a Program for Debugging](#)
- running, [Running a C++ Program](#)
- support, [GNU C++ Compiler](#)

C++11 (see GNU Compiler Collection)

c++filt

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

C11 (see GNU Compiler Collection)

Cachegrind

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

Callgrind

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

compatibility

- GNU Compiler Collection, [Language Compatibility](#), [Compatibility Changes](#), [Fortran 2003 Compatibility](#), [Fortran 2008 Compatibility](#), [Fortran 77 Compatibility](#), [ABI Compatibility](#), [Debugging Compatibility](#), [Other Compatibility](#)
- Red Hat Developer Toolset, [Compatibility](#)

compiling (see GNU Compiler Collection)

D

debugging (see GNU Debugger)

Developer Toolset (see Red Hat Developer Toolset)

documentation

- Red Hat Product Documentation, [Accessing Red Hat Product Documentation](#)

DRD

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

dwp

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

dwz

- documentation, [Additional Resources](#)
- installation, [Installing dwz](#)
- overview, [dwz](#)
- usage, [Using dwz](#)
- version, [About Red Hat Developer Toolset](#), [dwz](#)

Dyninst

- documentation, [Additional Resources](#)
- installation, [Installing Dyninst](#)
- overview, [Dyninst](#)
- usage, [Using Dyninst](#)
- version, [About Red Hat Developer Toolset](#), [Dyninst](#)

E

Eclipse

- configuration, [Using the Red Hat Enterprise Linux Toolchain](#)
- documentation, [Additional Resources](#)
- installation, [Installing Eclipse](#)
- overview, [Eclipse](#)
- usage, [Using Eclipse](#)
- version, [About Red Hat Developer Toolset](#), [Eclipse](#)

elfedit

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

elfutils

- documentation, [Additional Resources](#)
- installation, [Installing elfutils](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)
- version, [About Red Hat Developer Toolset](#), [elfutils](#)

eu-addr2line

- features, [Changes Since Red Hat Enterprise Linux 5.9](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-ar

- features, [Changes Since Red Hat Enterprise Linux 6.4](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-elfcmp

- features, [Changes Since Red Hat Enterprise Linux 5.9](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-elflint

- features, [Changes Since Red Hat Enterprise Linux 6.4](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-findtextrel

- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-make-debug-archive

- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-nm

- features, [Changes Since Red Hat Enterprise Linux 6.4](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-objdump

- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-ranlib

- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-readelf

- features, [Changes Since Red Hat Enterprise Linux 6.4](#), [Changes Since Red Hat Enterprise Linux 5.9](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-size

- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-strings

- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-strip

- features, [Changes Since Red Hat Enterprise Linux 6.4](#), [Changes Since Red Hat Enterprise Linux 5.9](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

eu-unstrip

- overview, [elfutils](#)
- usage, [Using elfutils](#)

F

feedback

- contact information for this manual, [We Need Feedback!](#)

Fortran programming language

- compiling, [Using the Fortran Compiler](#)
- running, [Running a Fortran Program](#)
- support, [GNU Fortran Compiler](#)

G

g++ (see GNU Compiler Collection)

GAS (see GNU assembler)

GCC (see GNU Compiler Collection)

gcc (see GNU Compiler Collection)

GDB (see GNU Debugger)

gfortran (see GNU Compiler Collection)

Git

- configuration, [Configuring Git](#)
- documentation, [Additional Resources](#)
- installation, [Installing Git](#)
- overview, [Git](#)
- usage, [Using Git](#)
- version, [About Red Hat Developer Toolset](#), [Git](#)

Global Support Services

- contacting, [Contacting Global Support Services](#)

GNU assembler

- documentation, [Additional Resources](#)
- installation, [Installing binutils](#)
- overview, [binutils](#)
- usage, [Using the GNU Assembler](#)

GNU Binutils (see binutils)

GNU Compiler Collection

- C support, [GNU C Compiler](#)
- C++ support, [GNU C++ Compiler](#)
- compatibility, [Language Compatibility](#), [Compatibility Changes](#), [Fortran 2003 Compatibility](#), [Fortran 2008 Compatibility](#), [Fortran 77 Compatibility](#), [ABI Compatibility](#), [Debugging Compatibility](#), [Other Compatibility](#)
- documentation, [Additional Resources](#)
- features, [Main Features](#), [Status and Features](#), [New Features](#), [Fortran 2003 Features](#), [Fortran 2008 Features](#)
- Fortran support, [GNU Fortran Compiler](#)
- installation, [Installing the C Compiler](#), [Installing the C++ Compiler](#), [Installing the Fortran Compiler](#)
- overview, [GNU Compiler Collection \(GCC\)](#)
- usage, [Using the C Compiler](#), [Using the C++ Compiler](#), [Using the Fortran Compiler](#), [Preparing a Program for Debugging](#)
- version, [About Red Hat Developer Toolset](#), [GNU Compiler Collection \(GCC\)](#)

GNU Debugger

- documentation, [Additional Resources](#)
- features, [Main Features](#)
- installation, [Installing the GNU Debugger](#)
- overview, [GNU Debugger \(GDB\)](#)
- preparation, [Preparing a Program for Debugging](#)
- usage, [Running the GNU Debugger](#), [Listing Source Code](#), [Setting Breakpoints](#), [Starting Execution](#), [Displaying Current Values](#), [Continuing Execution](#)
- version, [About Red Hat Developer Toolset](#), [GNU Debugger \(GDB\)](#)

GNU linker

- documentation, [Additional Resources](#)
- installation, [Installing binutils](#)
- overview, [binutils](#)
- usage, [Using the GNU Linker](#)

gprof

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

GSS (see Global Support Services)

H

Helgrind

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

help

- getting help, [Do You Need Help?](#)
- Global Support Services, [Contacting Global Support Services](#)
- Red Hat Product Documentation, [Accessing Red Hat Product Documentation](#)

L

ld (see GNU linker)

linking (see GNU linker)

M

Massif

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

Memcheck

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

memstomp

- documentation, [Additional Resources](#)
- installation, [Installing memstomp](#)
- overview, [memstomp](#)
- usage, [Using memstomp](#)
- version, [About Red Hat Developer Toolset](#)

N

nm

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

O

objcopy

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

objdump

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

opannotate

- overview, [OProfile](#)
- usage, [Using OProfile](#)

oparchive

- overview, [OProfile](#)
- usage, [Using OProfile](#)

opcontrol

- overview, [OProfile](#)
- usage, [Using OProfile](#)

opgprof

- overview, [OProfile](#)
- usage, [Using OProfile](#)

ophelp

- overview, [OProfile](#)
- usage, [Using OProfile](#)

opimport

- overview, [OProfile](#)
- usage, [Using OProfile](#)

opjitconv

- overview, [OProfile](#)
- usage, [Using OProfile](#)

opreport

- overview, [OProfile](#)
- usage, [Using OProfile](#)

OProfile

- documentation, [Additional Resources](#)
- installation, [Installing OProfile](#)
- overview, [OProfile](#)
- usage, [Using OProfile](#)
- version, [About Red Hat Developer Toolset](#), [OProfile](#)

oprofiled

- overview, [OProfile](#)
- usage, [Using OProfile](#)

R

ranlib

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

readelf

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

Red Hat Developer Toolset

- compatibility, [Compatibility](#)
- documentation, [Additional Resources](#), [Accessing Red Hat Product Documentation](#)
- features, [Main Features](#)
- installation, [Installing Red Hat Developer Toolset](#)
- overview, [About Red Hat Developer Toolset](#)
- subscription, [Getting Access to Red Hat Developer Toolset](#)
- support, [About Red Hat Developer Toolset](#)
- uninstallation, [Uninstalling Red Hat Developer Toolset](#)
- update, [Updating Red Hat Developer Toolset](#)

Red Hat Enterprise Linux

- documentation, [Additional Resources](#), [Accessing Red Hat Product Documentation](#)
- supported versions, [Compatibility](#)

Red Hat Subscription Management

- subscription, [Using Red Hat Subscription Management](#)

RHN Classic

- subscription, [Using RHN Classic](#)

S

scl (see Software Collections)**size**

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

Software Collections

- documentation, [Additional Resources](#), [Accessing Red Hat Product Documentation](#)

- overview, [About Red Hat Developer Toolset](#)

stap

- overview, [SystemTap](#)
- usage, [Using SystemTap](#), [Using Dyninst with SystemTap](#)

stap-merge

- overview, [SystemTap](#)
- usage, [Using SystemTap](#)

stap-prep

- overview, [SystemTap](#)
- usage, [Installing SystemTap](#)

stap-report

- overview, [SystemTap](#)
- usage, [Using SystemTap](#)

staprun

- overview, [SystemTap](#)
- usage, [Using SystemTap](#)

stapsh

- overview, [SystemTap](#)
- usage, [Using SystemTap](#)

strace

- documentation, [Additional Resources](#)
- installation, [Installing strace](#)
- overview, [strace](#)
- usage, [Using strace](#)
- version, [About Red Hat Developer Toolset](#), [strace](#)

strings

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

strip

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

support

- Red Hat Developer Toolset, [About Red Hat Developer Toolset](#)

SystemTap

- documentation, [Additional Resources](#)
- installation, [Installing SystemTap](#)
- overview, [SystemTap](#)
- usage, [Using SystemTap](#), [Using Dyninst with SystemTap](#)
- version, [About Red Hat Developer Toolset](#), [SystemTap](#)

V

Valgrind

- documentation, [Additional Resources](#)
- installation, [Installing Valgrind](#)
- overview, [Valgrind](#)
- usage, [Using Valgrind](#)
- version, [About Red Hat Developer Toolset](#), [Valgrind](#)

version

- version, [memstomp](#)